_____

AUDACITY SOUND EDITOR SOFTWARE

A USER GUIDE FOR AUDIO-VISUAL WORKERS

Prepared by Peter Appleton

Copyright 2008

All illustrations in this guide were created using Audacity v1.2.6

CONTENTS

_____

1. Introduction
   1.1. Purpose of this manual
      This manual is not intended to be a comprehensive guide to using Audacity as a sound editor program. Rather, it concentrates on those aspects of Audacity that are likely to be of most value to someone preparing a soundtrack for an AV sequence. It therefore concentrates on those aspects that apply to the background music and the voice-over, and the mixing of them both into the final soundtrack.

   1.2. Audacity's Help File
      It is not the intention of this document to replace the Audacity Help File. Please do take time out to read the contents of the Help File before reading this User Guide.

   1.3. Audacity's Reference Manual and Tutorials
      It is also not the intention of this document to replace the excellent Reference manual and Tutorials that can be found on the web at:
      http://audacity.sourceforge.net/manual-1.2/index.html

      This web source provides excellent explanations of:
         Sample Rate as expressed in Hertz (Hz)
         Sample Format as expressed in Bits
         Sound level as expressed in decibels (dBs)
         File formats including: AUP, WAV, MP3 and others
         The component parts of the Audacity Window
         Basic editing techniques including: Selection, Cut, Copy, Paste

      Again, please do take time out to explore this excellent resource before reading this User Guide.

   1.4. Audacity Wiki
      Also available via the Audacity web site http://audacity.sourceforge.net/ is a Wiki containing many articles of interest and value to anyone venturing into the complex world of sound-recording and sound-engineering.

   1.5. Audacity Forum
      The Audacity community also has its own Forum, accessible via the Audacity website. This, too, is worth some time spent on exploration.
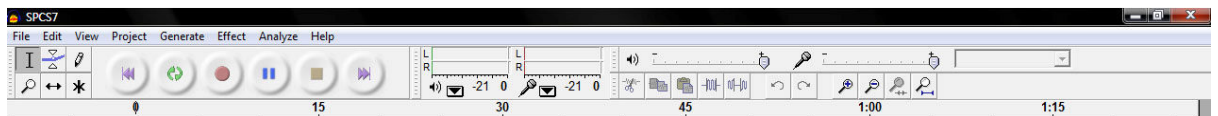
_____

2. Concepts and Terminology

    2.1. Audacity window – its component parts

        To ensure that we have a common vocabulary, let's take a brief look at the components of the Audacity window. This subject is covered in greater depth in the Audacity Reference Manual (see reference in Section 1.3). Here is a typical Audacity window.
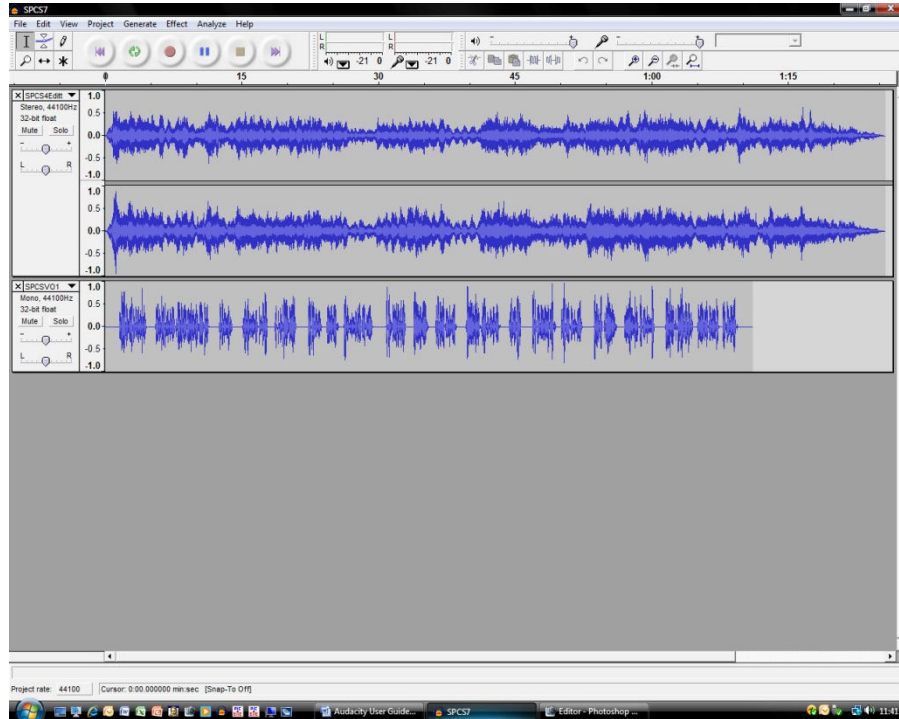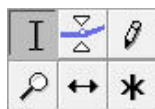

Let's look at the various parts of it:

At the top we have the Windows title bar, the Command bar, then various of Audacity's features and below them all, the timeline. In this example the timeline has tick-marks every five seconds.

Let's explore the Audacity features:

 These are the icons for Audacity's tools. Reading left to right, top to bottom they are: the Selection tool, the Envelope tool, the Draw tool, the Zoom tool, the Time Shift tool and the Multi-Tool Mode. Most of what we want to achieve will be done using just the Selection, Envelope and Time-Shift tools.

 These Playback and Record controls are intuitive to use.
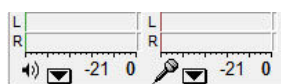
 The Level meters show the signal levels during playback (the left meter – shows green bars) and recording (the right meter – shows red bars).

_____

_____

The volume controls allow you to set a volume level on the playback (left slider) and record (right slider).

The Edit tools give you quick access to commonly used functions when editing your tracks. From left to right they are:

Cut	Removes the selected passage and places it on the clipboard
Copy	Copies the selected passage to the clipboard
Paste	Pastes the clipboard contents at the current cursor location
Silence outside	Replaces everything except the selected passage with silence
Silence selection	Replaces the selected passage with silence
Undo	Undoes the last action
Redo	Re-does the last undo
Zoom In	Increases the granularity of the timeline
Zoom Out	Let's you see the bigger picture, as it were
Fit selection in window	Adjusts zoom to do what it says
Fit project in window	Adjusts zoom to do what it says

2.2. Stereo and Mono recordings

The sample window we used in Section 2.1 showed two tracks: one stereo the other mono.
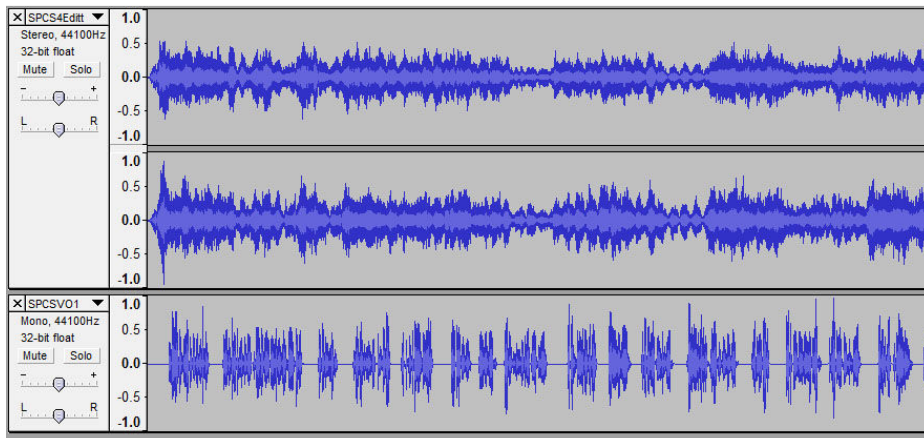
2.2.1. Stereo

A stereo track is one having separate left and right channels, the left channel being fed to the left loudspeaker and the right channel being fed to the right loudspeaker. In this context, left and right are as perceived by the listener sat in front of the loudspeakers.

2.2.2. Mono

A mono track is one having a single channel which, on playback, will be fed equally to the left and right loudspeakers.

_____

2.3. The Track Panel

The track panel at the left of each track contains the following features:
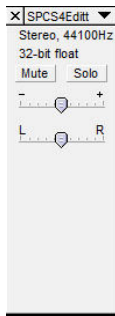
At the top is the track title and to its left, the "x" button which can be used to close this track (i.e. remove it from the window).

Immediately below the track title is the indication as to whether this is a Stereo or Mono track together with the Sample Rate (in this case 44100Hz) and the Sample Format (in this case 32-bit float). For detailed explanations of Sample Format and Sample Rate, see the Audacity Reference Manual mentioned in Section 1.3.

Beneath this information are the Mute and Solo buttons. The Mute button is used to prevent this track playing back. The Solo button is used to play back only this track. Note the you can Mute several tracks at the same time and you can Solo several tracks at the same time. In this latter case, you are effectively picking which tracks you want to hear and then having them play together.

Beneath the buttons is the Gain control. This allows you to have the track play back at a lower volume than that at which it was recorded or at a greater volume to that at which it was recorded. Below the Gain control is the Pan control. This controls how the signal is distributed between the left and right channels.

3.   Importing and ordering the soundtrack components

3.1. Importing the soundtrack components

It is assumed that all your soundtrack material: music files, voice-over, actuality recordings, etc. already exist as WAV files. To bring these into an Audacity project use the Project…Import Audio… command.
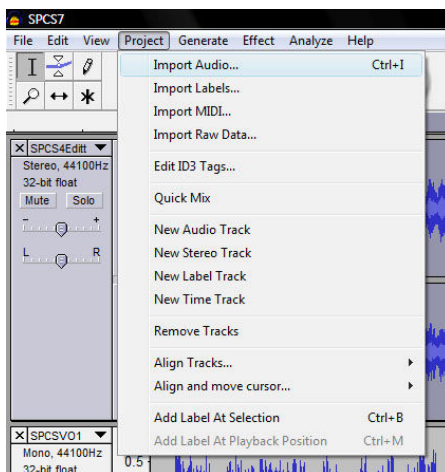
This will give you a standard file navigation window that you can used to locate and identify the sound file that you want Audacity to import into this project. As each one is imported it will appear as a separate track in the Audacity window.

As you can see, this command also allows you to add a new, empty stereo track or a new, empty mono track (called audio track in the list of options).

_____

3.2. Ordering the soundtrack components

You can change the order in which the tracks are displayed from the track options dropdown.
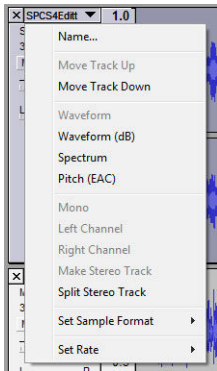
To do this you first need to click on the track title to get the drop down and then click on Move Up or Move Down. An alternative technique is simply to left click anywhere in the Track Panel and then drag the track up or down to the required position.

4. Creating new tracks

4.1. Creating a new mono track

As we saw in section 3.1, we can create a new, empty mono track by using the Project…New Audio Track command

4.2. Creating a new stereo track

As we saw in section 3.1, we can create a new, empty stereo track by using the Project…New Stereo Track command.

4.3. Creating tracks by splitting stereo tracks

If your voice-over recording is a stereo track, it might be worthwhile splitting it and converting one channel to mono and deleting the other channel. This will ensure that your voice is placed centrally on the stereo sound stage in the final mix down. To do this, proceed as follows:
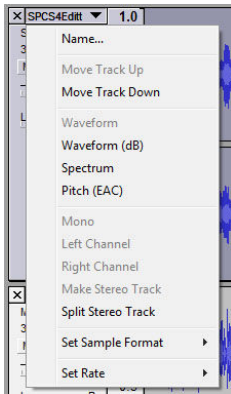
- Select your voice-over track by clicking on the Track Panel
- Call up the track options by clicking on the Track Title
- Click on Split Stereo Track
- Select one of the tracks thus produced
- From the track options, click on Mono
- Select the other track of the pair
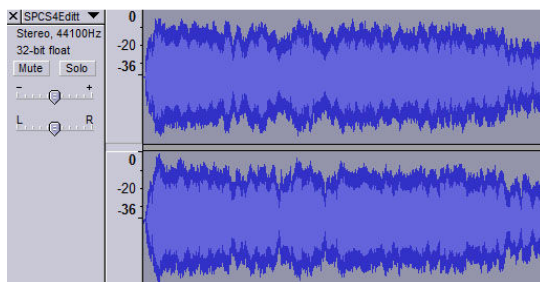- Click on the "x" button to remove it from this project

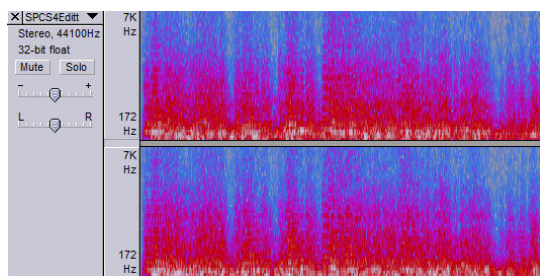5. Display of tracks
    5.1. Display of tracks
        There are four possible display types shown in the track options drop-down: Waveform, Waveform (dB), Spectrum and Pitch (EAC). A sample of each is shown below.
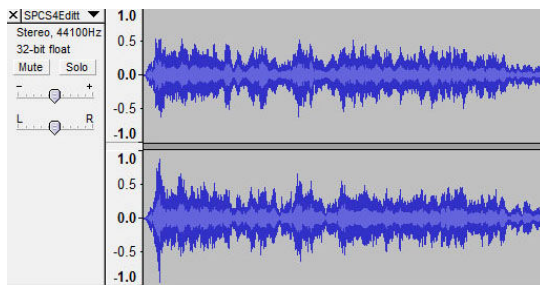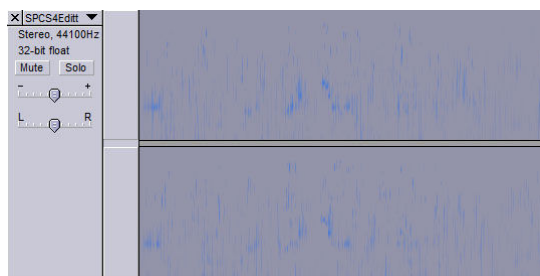

Waveform

Waveform(dB)

Spectrum

Pitch(EAC)

For most usage the Waveform will be the best choice. However, when doing a final mix-down using the Envelope Tool (see later) it can be advantageous to change to the Waveform(dB) display. I have not yet found a need to use either the Spectrum or Pitch(EAC) display forms.

_____

6. Noise Reduction
   6.1. Setting expectations
   Before we dive into the "how to do it" bit let's get one thing clear: you cannot make a poor recording better by using Effect…Noise Removal. (In my opinion Audacity has picked a poor term for this. You cannot remove noise. What you can do, sometimes, is reduce the noise). All that Effect…Noise Removal will do is make a poor recording even worse! What you can do is use Effect…Noise Removal to make an already good recording even better - sometimes.

   In this respect sound file manipulation is very different to image file manipulation. You can use Photoshop to remove unwanted objects (noise) from your images. In a sound recording, the frequencies involved in the noise (the frequencies that you want to remove) are the same frequencies that will be present in your voice or the music. You cannot remove the noise frequencies without having an effect (usually both detrimental and audible) on the rest of the track.

   Just as with your photography where, I hope, you try and get the best possible result at the time of taking, so with your sound recordings. You should strive to obtain the best possible original recording. Remember, the best form of noise reduction is not to record it in the first place.

   6.2. Concepts of Noise Reduction
   It helps if you understand the concepts behind noise removal filters. You give the filter a piece of recording that contains just the noise so that it can analyse it and understand what noise looks like digitally-speaking. Then you give it the whole or part of the recording and ask it to remove the noise. In broad terms, all it does is try and take out the given frequencies where they are the only thing present and then maybe adjust the volume elsewhere by reference to the amount of volume present in the noise to try and mask its presence.
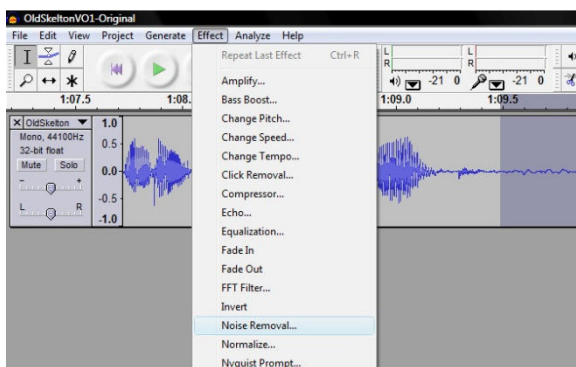
   So, with that out of the way, let's look at how you can use Audacity's Noise Removal tool.

   6.3. Using Effect…Noise Removal
   Begin by selecting a piece of your recording that has just the noise present. See example below:
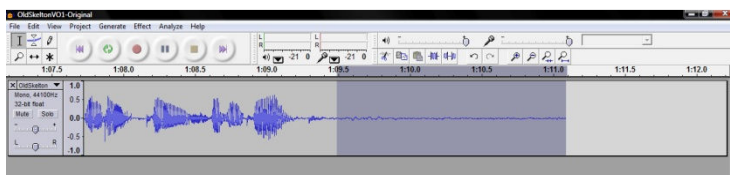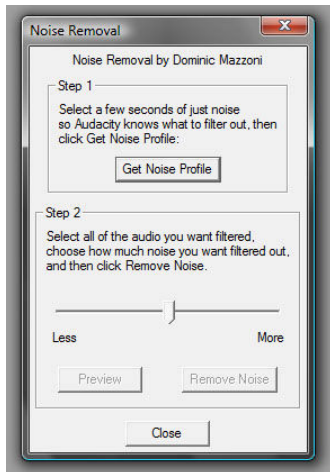

   Then use the Effect…Noise Removal… command.

_____

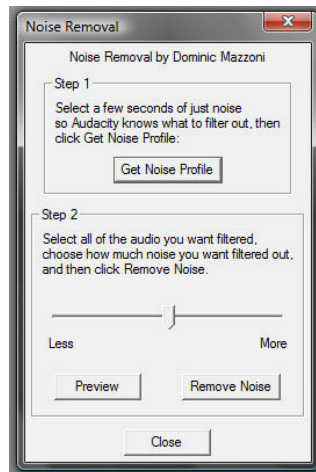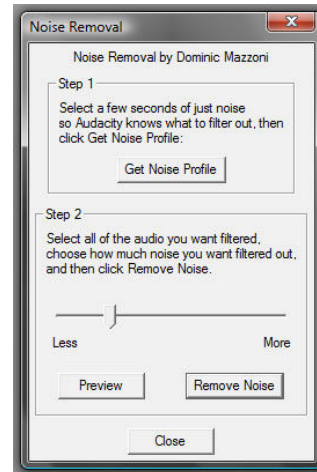| Figure 1 | Figure 2 | Figure 3 |

In the pop up window (Figure 1), click on the Get Noise Profile button.

 Now select that part of the recording on which you want to try and reduce the noise and use the Effect… Noise Removal… command again (Figure 2). Note that the Preview and Remove Noise buttons are no longer greyed out. At this point you need to experiment with different amounts of noise reduction. You do this by using the slider and the Preview button. Once you've found a level that gives the best amount of reduction with the least amount of introduced distortion (Figure 3 is the setting I find generally works for me), you can click on the Remove Noise button.

I stress once more, there is no substitute for getting the quality right in the original recording. Noise Removal is not a "magic wand" that will turn a poor recording into a good one. It won't even turn it into a passable one.


7. Manipulating tracks
   7.1. Adding silence to a recording
        There will be occasions when you would like some silence at a particular point in your soundtrack. I like to start and end all my sequences with a black slide onscreen for about 4-5 seconds. So I often want to add a few seconds of silence to the start and end of my soundtrack. Let's look at this technique:
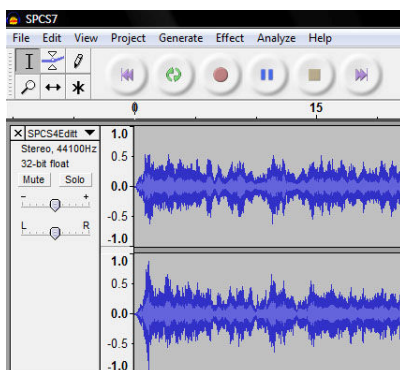
Here's my music track loaded into Audacity. Note that the Selection tool is the active tool.

_____
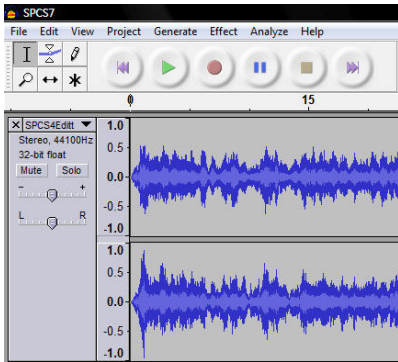
Next we must tell Audacity where we want the silence inserting. I want it at the beginning of the track. So here, I have just positioned the cursor at time zero by left clicking just inside the left edge of the waveform area.
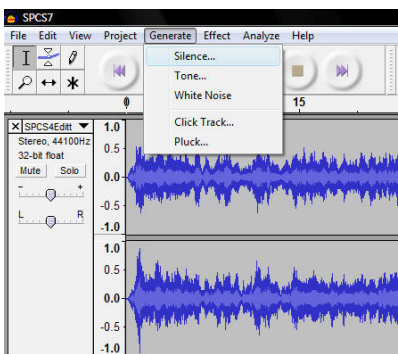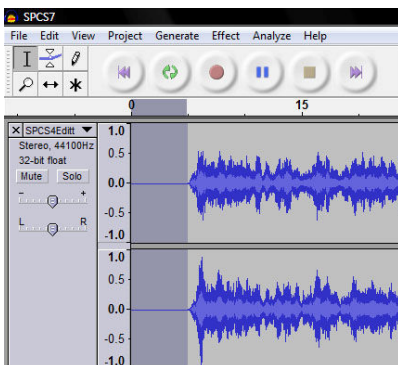
Now we can use the Generate… command.

And select the Silence… option.
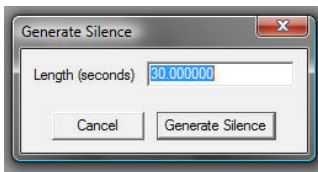
In the pop-up we can specify the number of seconds of silence that we want (30 is the default). In our worked example I would key 5 in this pop-up. Press the Generate Silence button to add the silence.

And here's the result. Our music track now begins with 5 seconds of silence.

7.2. Changing the length of a recording

When building an AV sequence, have you ever arrived at the point where you've found an ideal piece of music but it is either a just a bit too short or just a bit too long?

Well, all is not lost, Audacity enables you to lengthen or shorten a piece. You do it by using the Effect…Change Tempo… command. Let's take a look at this technique.
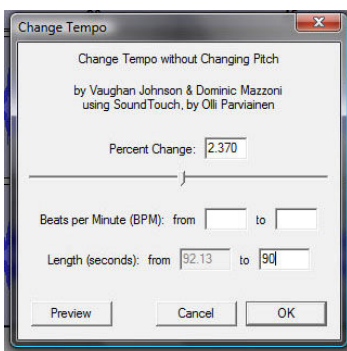
Here's a piece of music whose length we'd like to change.
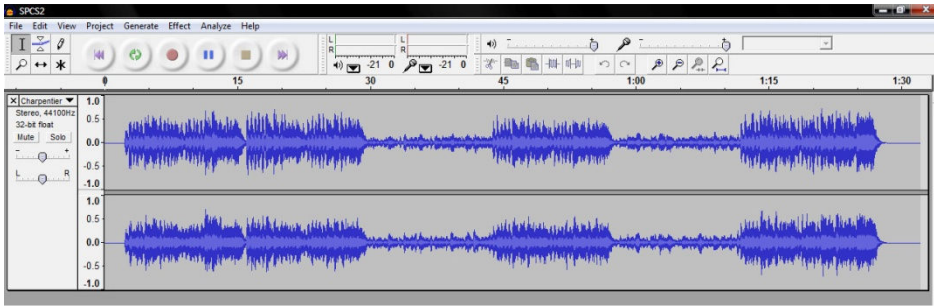
First we need to select the track whose duration we wish to change. Then we use the Effect… Change Tempo… command.
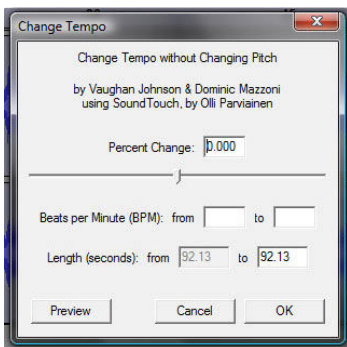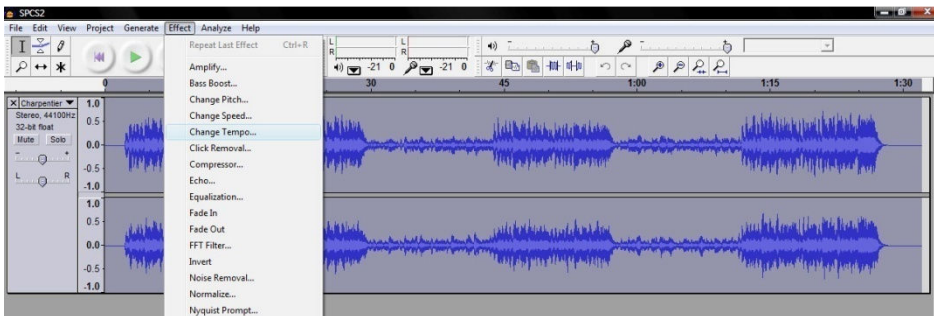
In the pop-up window, we note the duration of the piece (92.13 seconds) and over-key this value with whatever we want it to be.

In this example I have set it up to give a final length of 90 seconds.

I find that changes up to about plus/minus five percent variation are almost undetectable. Changes between five percent and ten percent are usually detectable but sometimes not intrusively so.

As with all things in AV – be prepared to experiment. And remember, your audience sees and hears only the finished product. You alone know how it all started out.

_____

7.3. Applying fades

There are three ways of applying fade to a track: Effect…Fade In, Effect…Fade Out and the Envelope Tool. This section deals with the first two, the Envelope Tool has a section all to itself later on.

For music that has been ripped from audio CD, the ideal situation is to use the "natural" fades that exist in the original recording. If, however, you feel that your best option is to cut part of the music track and apply fades to it then here is what you do.

Here's a piece of music that has had its final part cut away.
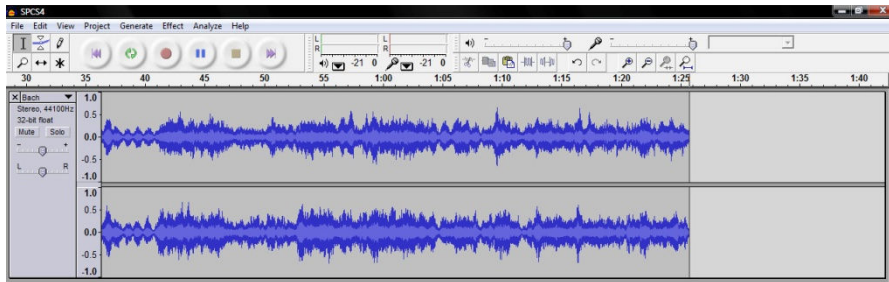
We want to apply a fade out to the end of the piece. Firstly, we must decide how long a fade to have. A fade in can, typically, be quite short; maybe only 1or 2 seconds. A fade out, typically, needs to be longer; maybe 5 to 10 seconds or even more. You can use the Gain slider on the track panel during playback to practice the fade and assess the overall effect before committing to apply the fade. But do remember to return the Gain to its original position.
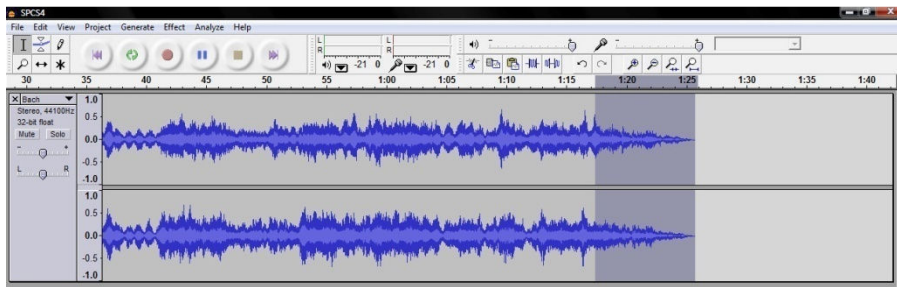
We must now select the part of the music to which we want to apply the fade. The simplest way is to click just off the end of the track and then drag back to the left.
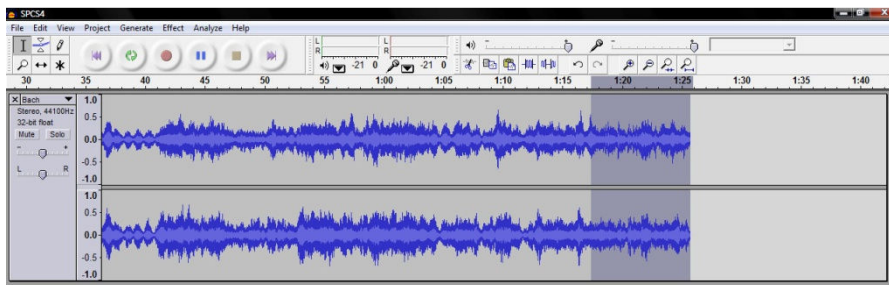
Now all we need to do is use the Effect…Fade Out command.

_____
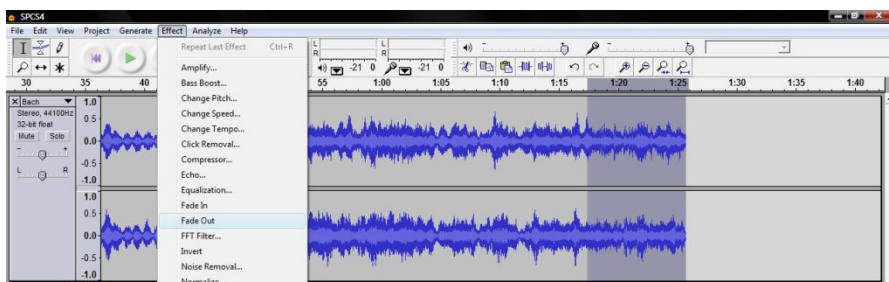
And there is our fade out applied.

Obviously, for a fade in you would be working at the other end of the track, but the basic technique is the same.
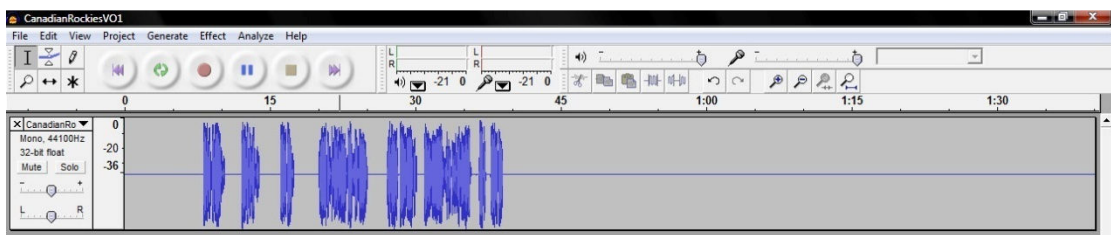
7.4. Using Effect…Normalize…

When making an original open-microphone recording we must ensure that we have controlled the input signal level to avoid "clipping". Clipping occurs when the signal strength exceeds the capability of the recording equipment to capture it without distortion. It can be best understood by loading a bad recording into Audacity and changing the view to show the Waveform(db) display.
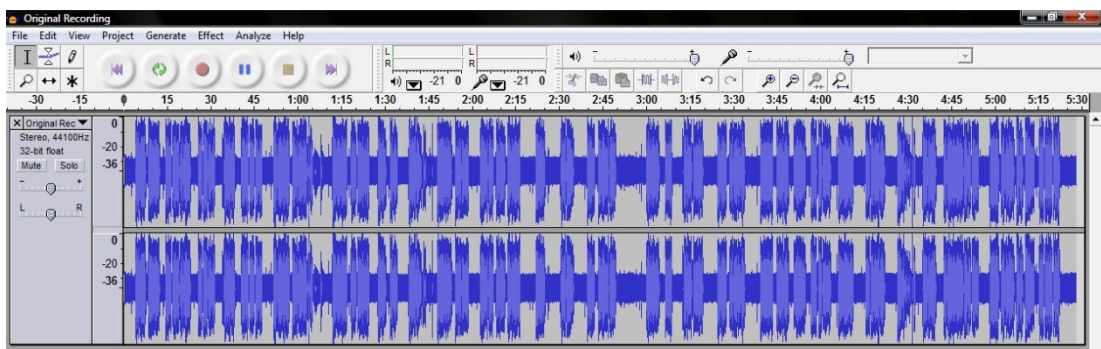

Above can be seen a badly clipped recording. Many of the spikes extend all the way to the edge of the track window. A really bad recording will have "flat spots" where the sound level is constantly hard up against the edge of the track window. (This recording has many other faults but it's the clipping we're interested in). This recording is so badly clipped that nothing can be salvaged from it. A re-take is required with recording levels kept under better control.

Commercial recordings should have been properly engineered and should not have any clipping. Indeed, they should appear to have their peak levels around the -3dB level on the Waveform(dB) display.

If your voice-over recording has just occasional and very short duration clipping, you can sometimes achieve a cleaner result by normalizing your recording.
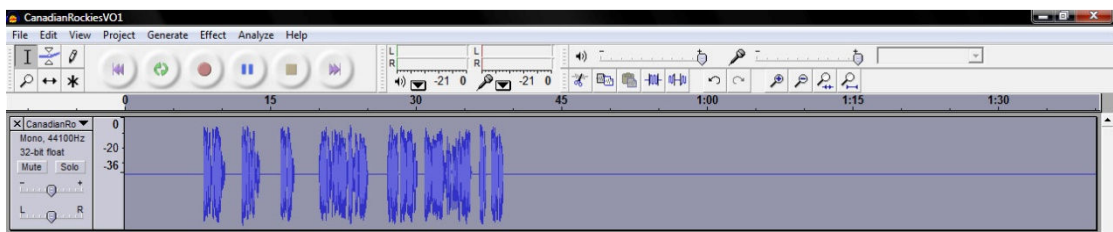
The previous illustration shows a voice-over with just very occasional, slight clipping. This can be salvaged by normalization. To do this you simply select the track to be normalized and then use the Effect…Normalize… command.
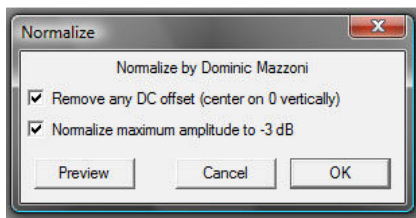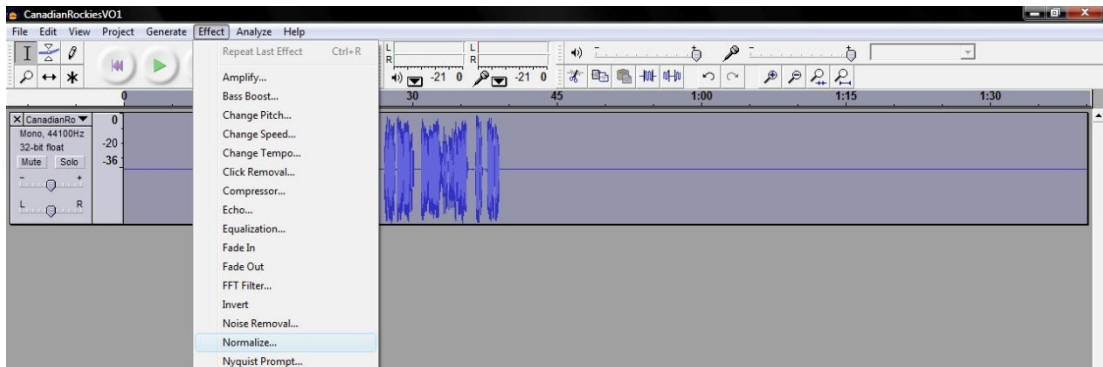
 I always take the default values in the pop-up window.

And here's the result of normalization:

8. Mixing tracks
8.1. Simple mixing
Let's assume that you have two tracks loaded into Audacity: a music track and a voice-over track. And let's also assume that you've done all the processing you want on them and are ready to do the final mix down. Finally let's assume that all or nearly all of your images have voice-over.

_____
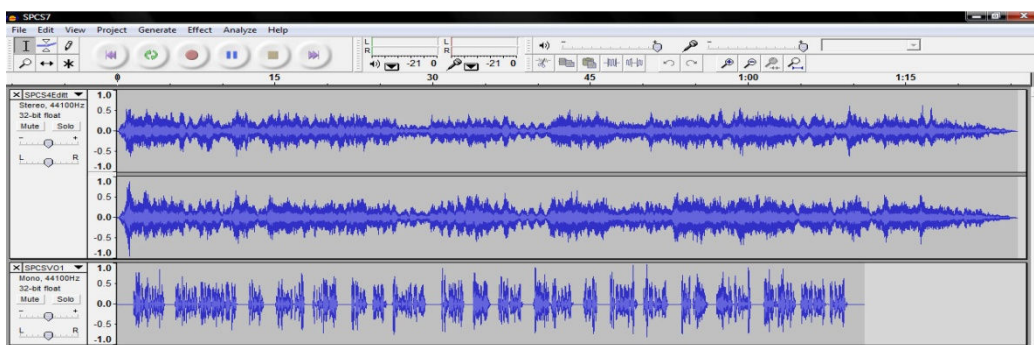
In this situation all you need to do is hold back the music volume all the way through the final mix. This is most easily achieved by lowering the gain on the music track to -18dB or even -21dB depending on the overall volume of the music track.
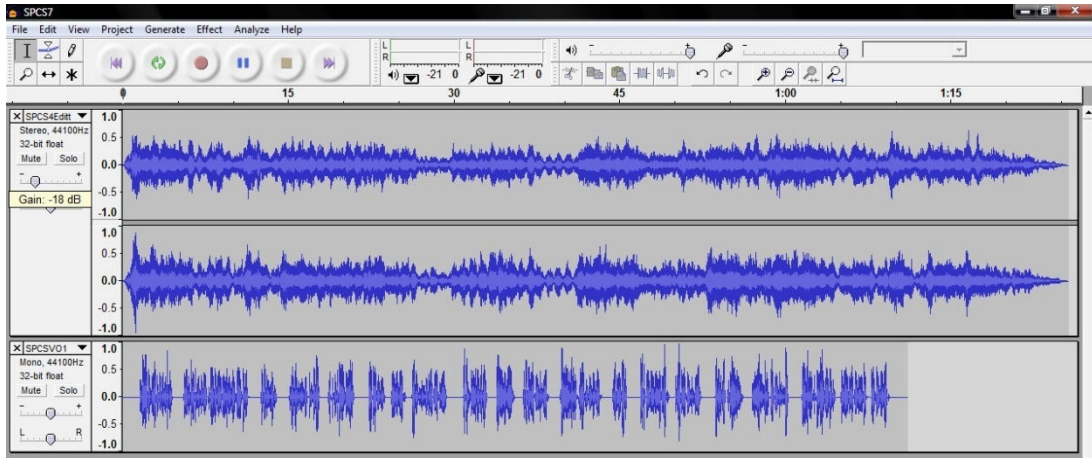

You can preview this mix simply by changing the gain and then playing back. Note that the waveform display of the music track does not change appearance to show the reduced gain level.

If you are happy with this mix then export the entire project as an MP3 file to use in your PTE project.

8.2. Complex mixing

If your voice-over applies only to some images, with short runs of images having no voice-over, you have a more complex mix to do. You will still want the music lowering to -18dB or even -21dB where it conflicts with your voice-over. But on the parts where you have just slides and music you will probably like to have the music a little louder.

This degree of control is possible using the Envelope Tool. Here's how to set about using this tool. Below we have a mixing set-up involving a stereo music track and a mono voice-over track. The music track has been selected.

_____
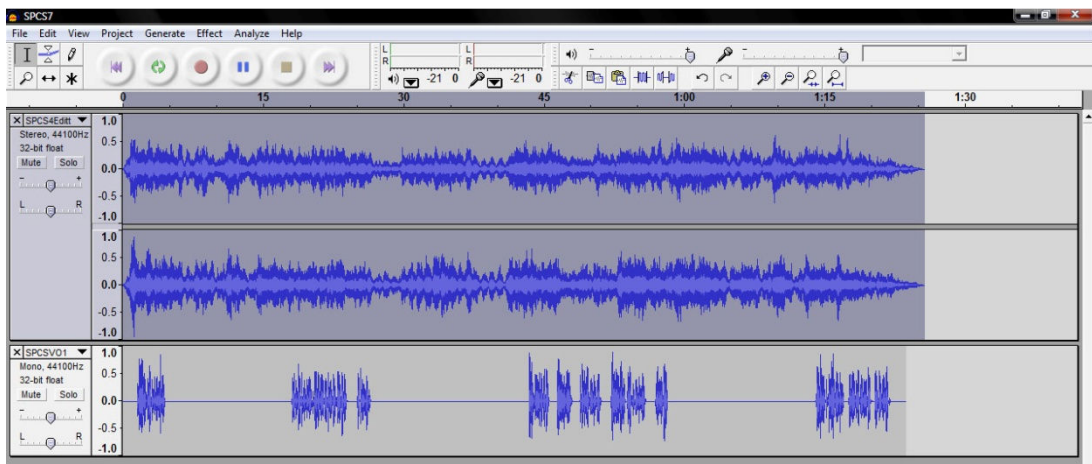
The next step is to choose the Envelope tool, giving this display:
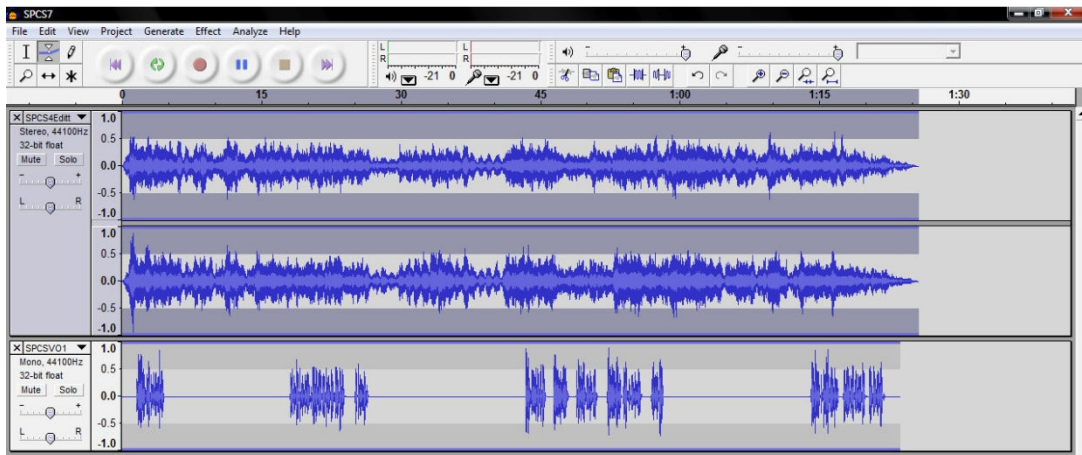
When using the Envelope tool we always want to see the timeline in some detail so that we can position our envelope points with reasonable accuracy. So here, I have zoomed in until the timeline is showing tick-marks every half-second.
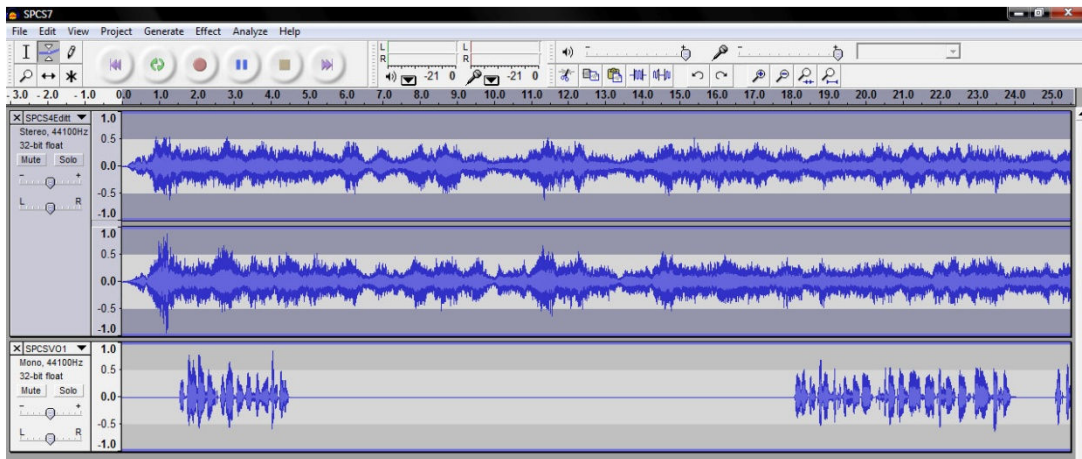

To get to the next sample I have done several things:
- I have changed the display of the stereo track to show the Waveform(dB) display
- I have positioned the mouse pointer on the blue line at the top of the stereo track, as close to time zero as I could, and have then dragged the blue line down to just above -20dB on the vertical scale
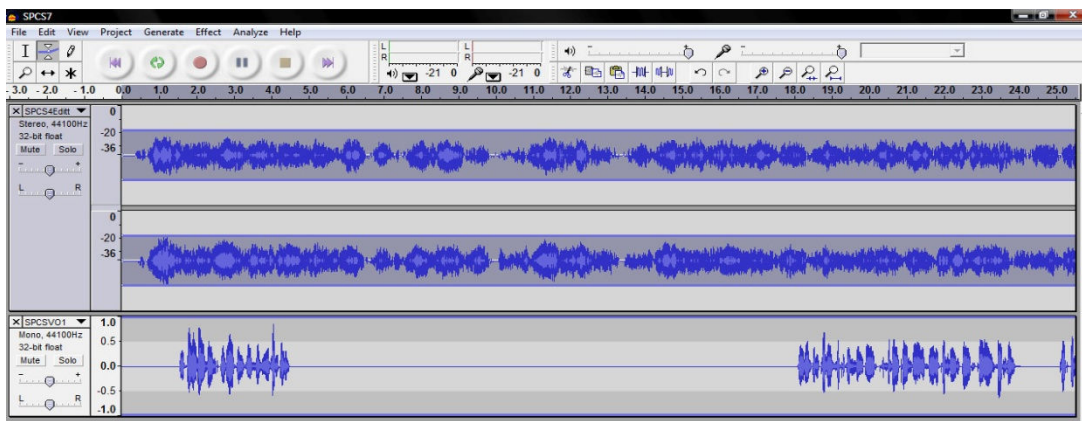
That has set the base level for the music so that it doesn't drown out the voice-over. It's the equivalent of lowering the gain to -18dB as we did in the simple mix in the previous section.

But we would like to lift the volume up a little during those parts of the sequence where we have no voice-over. So we need to place some more envelope points by clicking on the top blue line of the stereo track. We need one just as the piece of voice-over ends and another one about one second later. Then we need another pair at about one second before the next voice-over starts. We can then drag the appropriate points up a little so as to lift the volume of the music when there is no voice-over. The display takes on this kind of appearance after we have done that.
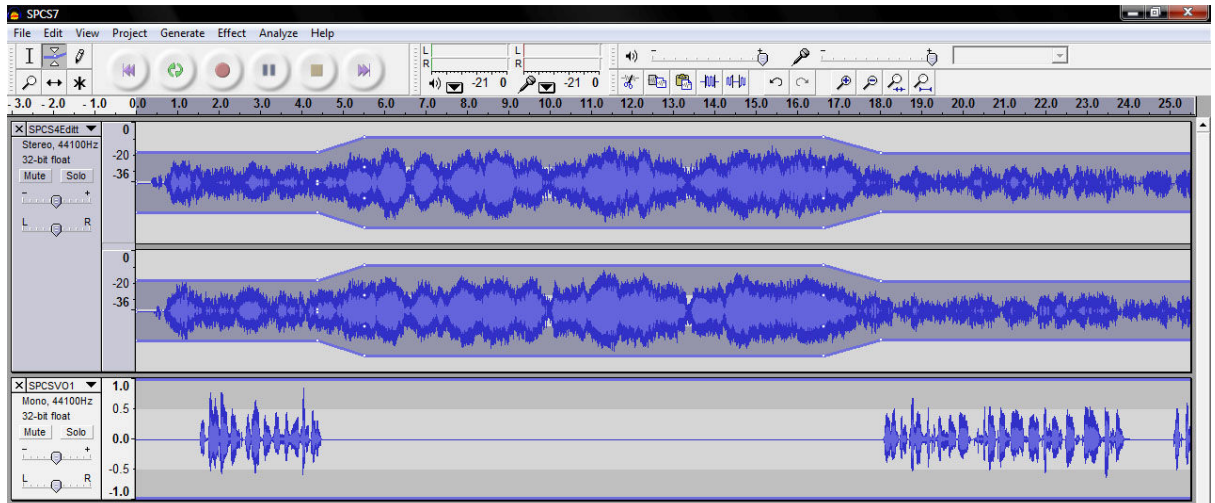
Artistically, you may find that you don't want to raise the music volume all the way back to full volume. Experiment and find out what works best for you.

Once you are happy with this mix then export the entire project as an MP3 file to use in your PTE project.