AudioSlicer User Guide

AudioSlicer programmed by Bernd Heller User guide written by Peter Watkins 2006-09-07 This user guide was written by Peter Watkins.

AudioSlicer was created by Bernd Heller. Copyright © 2005 Bernd Heller. All rights reserved.

AudioSlicer is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

AudioSlicer is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with AudioSlicer; if not, write to:

Free Software Foundation, Inc. 59 Temple Place, Suite 330 Boston, MA 02111-1307, USA

Contents

Introduction	5
System Requirements	5
Concepts	6
-iles	6
Audio Files	6
Split Files	6
Slice Files	6
Tracks and Slices	6
Silence	.7
Audio Content	7
Common Tasks	8
Load an Audio File into AudioSlicer	8
Set the Parameters for Slicing (Set the Preferences)	8
Save a Split File	8
Audition the Slices	8
Split the Slices	9
Export the Slices	.9
Reference 1	0
Main Menu 1	0
Application Menu 1	0
File 1	0
Edit 1	1
Slices1	1

Window1	12
Help 1	13
Preferences Window 1	13
Defaults for New Documents1	13
Silence Analysis 1	13
Export1	14
Split Document Window1	14
Toolbar1	14
Play Settings1	14
Track and Slice List1	15
Inspector Window 1	15
Overview1	15
Details1	15
Appendix A: Keyboard Shortcuts1	16
Application Menu1	16
File Menu1	16
Edit Menu1	17
Slices Menu 1	17
Window Menu1	17
Help Menu 1	17
Appendix B: FAQ1	18
Q: Why can't I use system software prior to OS 10.3?1	18
Q: Can I add more than one MP3 file to my split file?1	18
Appendix C: Troubleshooting 1	19

Introduction

Thanks for trying AudioSlicer. This user guide will show you how AudioSlicer works, walk you through some common tasks, and provide a handy reference.

AudioSlicer is an application that finds silences in an audio file, allows you to split the file into several smaller ones, and helps you name or tag those files properly. Currently only MP3 files are supported but other audio formats may be added in the future.

Most other tools do this split in one shot according to criteria you set beforehand. Once you've run your file through the tool's process, you're stuck with the divisions it chose. AudioSlicer finds all the silences in an audio file and lets you configure how you want to split the file at each silence. You can audition the audio around each silence point (i.e., listen to the audio from 3 seconds before the silence point through 4 seconds after). Then you can decide if you want to split the audio file at that point.

Splitting an audio file with other tools is similar to sculpting with granite—once you've chipped off a piece, you can't reattach it. Working with AudioSlicer is similar to sculpting with clay—you can detach and reattach a lump as many times as you like. It's this flexible method of auditioning possible split points that makes AudioSlicer much more powerful than the trial and error method other tools employ. In addition, splitting MP3 files is loss-less—AudioSlicer does not decode and re-encode the file.

System Requirements

- Apple Power Macintosh Computer
- Mac OS X 10.3 (Panther) or later

Concepts

Files

In AudioSlicer, you'll encounter audio files, split files, and slice files. Figure 1.1 shows how a split file references an audio file and splits it into slices.

Audio Files

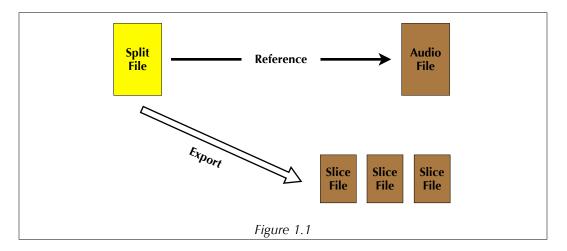
An audio file is the file you want to split into pieces. You'll be working with one audio file at a time since one split file is referenced to exactly one audio file. No matter what you do in AudioSlicer, your audio file will remain unchanged.

Split Files

You may think of a split file as your project file. It holds information about where you want to split the audio file. Whenever you work with AudioSlicer, you're editing your split file. When you save an AudioSlicer document, you're saving a split file. The split file does not contain the audio file, it references to it. For this reason, split files don't take much space on your disk.

Slice Files

A slice file is a chunk of audio extracted from from the audio file. Usually you'll export many slice files from a single audio file. Slice files are AudioSlicer's final product.

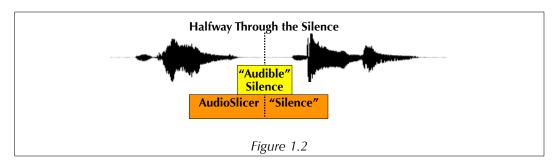


Tracks and Slices

Imagine slices as the DNA for slice files. Each split file can contain many slices, and when you export the slices from a split file, you get slice files. Slices are also referred to as tracks in AudioSlicer. The split file window looks similar to a playlist in a music player, and slices are similar to tracks in a music player.

Silence

We usually think of silence as the absence of noise. I'll call this "audible silence." In AudioSlicer, however, you'll find regions of "silence" that actually contain some audio content (noise). To make these regions of silence AudioSlicer finds some audible silence, and it picks a point halfway through that audible silence. When you audition (listen to) a file in AudioSlicer you can tell it to play a region of audio on both sides of this halfway point. This region of audio is labeled "AudioSlicer Silence" in figure 1.2. The region of audio where there isn't any noise is labeled "Audible Silence." By default, AudioSlicer will play a beep when it reaches the halfway point in an AudioSlicer Silence.

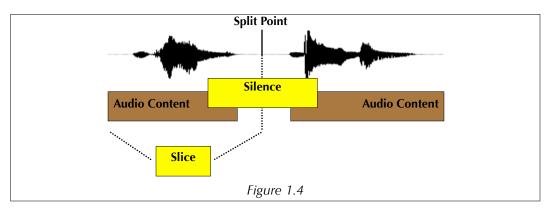


Now look at figure 1.3. It's a portion of the split file window in AudioSlicer. Here is the same AudioSlicer silence with audio content on either side. In the future, I'll refer to regions of silence like those in figure 1.3 as merely "silence" and everything else as "audible silence."

Title			Start Time
▼400 More Words			0:00'00.0
Audio Content		►	0:00'00.0
Silence		►	0:00'01.1
Audio Content		►	0:00'01.7
Cilonea	11	b.	0.00/01 0

Audio Content

Audio content is anything that isn't silent. You can turn any piece of audio content into a slice. Figure 1.4 shows how audio content turns into a slice. So if slices are baby slice files—pieces of audio content are baby slices.



There is a subtle difference between audio content in the split file window and slices. When you play a piece of audio content, only the audible noise is played. A slice, however, encompasses all the audio from one split point to another. This means that once you make a slice, it could contain some silence as well as audio content.

Common Tasks

Load an Audio File into AudioSlicer

Before you can begin working with AudioSlicer, you must load an audio file. Choose Open... from the file menu and navigate to the MP3 you want to load. AudioSlicer is nondestructive. It loads a copy of the MP3 file into memory. It won't split your original file into pieces. Once the file has finished loading, it will appear as a track in the split file window (the split file window's title bar will read "Untitled" at this point).

Set the Parameters for Slicing (Set the Preferences)

To successfully split a file, you'll need to fine tune settings in the preferences window. Choose Preferences... from the application menu.

The values in the Defaults for New Documents section control what you hear when you audition each silence. The default values of 2 seconds for both play intervals work well for me. By default, AudioSlicer will beep when it plays over each split point. If this annoys you, you can uncheck Beep in Middle of Silence.

The values in the Silence Analysis section are very important. The minimum silence duration will depend on what type of audio you're splitting. If you're splitting songs in an album apart, you'll want a higher value—around 1 to 3 seconds. If you're trying to split spoken words or phrases apart, you'll need a value around one second or lower.

Your maximum silence volume will depend on how noisy your recording is. If you have a recording with a lot of background noise, choose a higher value. If your recording is clean, a low value will keep AudioSlicer from interpreting audio content as silence.

You can choose where you want your split point to occur in a section of silence. Suppose there is a four second silence between each piece of content in your audio file. You may want to put your split points near the end of the silence or at the beginning. If you put the split point at the beginning, the slice preceding the split point will have a short amount of silence at its end while the slice after the split point will have a longer silence at its beginning. You can't set the exact location of each split point in a silence individually. Instead you set the relative split point percentage for the entire split file. A higher percentage will put each split near the end of its silence while a lower percentage will put it near the beginning.

The preference you set apply to all new split files. So if you want them to take effect after you've loaded an audio file, you must reload the audio file.

Save a Split File

To save a split file, choose Save from the File menu. To save a copy choose Save As... Split files are not MP3 files. You can't load a split file into iTunes. You can export slices from a split file and load them into iTunes, however.

Audition the Slices

Once you've set your preferences and loaded an audio file, it's time to listen to your split file and chose where you want to split it. This is called auditioning. First listen to some of

the silences and audio content. When you hear the beep in the middle of a silence, is it in the right place? If the you're getting beeps where there shouldn't be any, try dragging the left silences range slider to the right. This will get rid of some of the shorter silences. Dragging the right silences range slider to the left will eliminate the longer silences.

Split the Slices

After auditioning, you can start splitting your silences. Each time you split a silence, the slice containing that silence will be divided in two. One way to split a silence is to select it and choose Split Slices from the Slices menu. Another way is to click the Split button to the right of the silence you want to split.

Export the Slices

To export the slices you've made, choose Export Splitted... from the File menu or click on the icon that looks like a stack of pages in the toolbar. A file save dialog box will pop up. Unlike other save dialog boxes, you don't specify a filename, but rather a filename format. A filename format consists of keys and any other text you want to include. Keys are surrounded by square brackets (i.e. [title] is a key). These keys are the ID3 tags for each track or slice. So if you specify a filename format of "[trackNumber] - [title]" AudioSlicer will save each slice in an individual MP3 file with a name consisting of the track number, a hyphen, and the title of the track.

Reference

Main Menu	
Application MenuAbout AudioSlicer	Shows a dialog box listing the creators of AudioSlicer. This is also a good place to check the version number.
• Preferences	Displays the preferences window. There you can set the default settings for new documents. Turn to Preferences Window later in this section to learn more.
• Services	The services in the services menu mostly operate on text. Since AudioSlicer is an audio program, you won't find much of interest here.
Hide AudioSlicer	Hides AudioSlicer.
• Hide Others	Hides every application except AudioSlicer.
• Show All	Shows all open applications.
Quit AudioSlicer	Terminates AudioSlicer. You'll be asked to save your open split files if you haven't already.
File	
• Open	Opens an audio file (i.e., an MP3 file) or a split file. You don't need to select a format when opening files, Au- dioSlicer will determine if the file is an audio file or a split file automatically.
Open Recent	This menu shows a list of split files you've worked on recently. It will not show imported audio files.
• Close	Closes a split file. You'll be asked to save if you haven't.
• Save	Saves a split file. If you want to save an audio file, choose Export Splitted instead.
• Save As	Allows you to save a copy of your split file with a new name.
• Revert	Loads your saved split file and discards any changes you made after you saved. This item is dimmed if you ha- ven't saved your file yet.

• Export Splitted	Once you've made all the splits you want, choose this menu option to export each slice (audio content be- tween the splits) as a separate MP3 file.
Edit	
• Undo	Retracts the last change you made to your split file.
• Redo	Retracts an undo. AudioSlicer has multiple levels of undo and redo.
• Cut	Copies the selected text to the clipboard and deletes it from the location it was cut from. The Cut, Copy, Paste, and Delete commands in AudioSlicer work only on text. You can't use them on segments or tracks at this time.
• Сору	Copies text without modifying the location you copied it from.
• Paste	Inserts text you copied or cut.
• Delete	Deletes selected text.
• Select All	Selects all the text in a field, or all the tracks and slices in your split file.
• Spelling	You probably won't be checking much spelling in Au- dioSlicer, but you could use the tools in this menu on the comments field for a track.
Slices	
• Split Slices	Once you've listened to a silence and are happy with the split point AudioSlicer has chosen, use this com- mand to split your track at the silence you have se- lected. This menu item will be dimmed if you select audio content since audio content doesn't have a split point. The Split Slices menu item does the same thing as the Split button to the right of a silence. When you split a slice, the single track the silence was in will turn into two tracks. The first track ends at the silence you split and the second begins there.
• Join Slices	Oops! You split a silence and you didn't mean to. Just select the silence you split and choose Join Slices. This menu item will be dimmed if you select audio content or silence that hasn't been split. The Join button to the right of each split silence that does the same thing as this menu item.
• Expand	Reveals all the silences and audio content contained within a track. You must have a collapsed track selected to use this command.

• Collapse	Hides the silences and audio content in a track. You can also expand and collapse a track by clicking the triangle next to it.
• Play Previous Silence	Plays the silence previous to the audio content or si- lence you currently have selected. If you have the first audio content or silence in your split file selected, this command does nothing (there is no previous silence to play).
• Play Next Silence	Plays the next silence in a track after the audio content or silence you currently have selected. Just like Play Previous Silence, this command doesn't do anything if you have the last audio content or silence in your file selected.
• Rename Serially	Renames your tracks by incrementing a number at the end of each track name. You must have more than one track selected to use this command. Once you choose Rename Serially a dialog box will pop up with two fields: Start Name and Padding Length.
	Start Name is the name you want the first track in your selection to have. For Rename Serially to do you any good, you must include a number at the end of your start name. Once you've clicked Rename, the second track in your selection will end with a number one higher than the first, the third track, one higher than the second, and so on.
	Padding Length tells AudioSlicer how many leading zeros to put in the number at the end of each name. So if you entered 3 in the Padding Length field, the numbers on the ends of your tracks would look like 001, 002, 003, etc.
Renumber Serially	Incrementally numbers the tracks starting with track number of the first track.
Show Info	Brings up the inspector palette. The inspector palette shows the info for whatever you currently have selected.
Window	
• Zoom	Expands the front window to fill the screen or contracts it back to its original size.
• Minimize	Minimizes the front window to the dock.
Bring All to Front	Moves all AudioSlicer windows in front of any other applications windows. This command won't move windows out of the dock, however.

Help

AudioSlicer Help	Displays help for AudioSlicer using the Apple Help Cen- ter. Currently there is no help center book for AudioSli- cer, but there will be one soon.
 About partnermine.com 	Displays the partnermine.com splash screen again. From the splash screen, you can go to the partnermine.com website. Partnermine.com was started by Bernd Heller and Tim Trent. It's the dating site that does it right.

Preferences Window

You can set defaults for new documents in the preferences window. Some settings in the preferences window also appear in the split document window.

Defaults for New Documents

When you're auditioning the silences in a split file, it helps to listen to some audio before and after the split point. How much audio you want to listen to will depend on your personal tastes. After working with a few split files you'll probably find some default values here that suit you.

Play Interval Before Silence	The number of seconds of audio to play before each split point.
Play Interval After Silence	The number of seconds of audio to play after each split point.
• Beep in Middle of Silence	AudioSlicer can beep when the split point oc- curs. This gives you an audible mark where si- lence will be split. You might find this beep an- noying, however.

Silence Analysis

AudioSlicer analysis the regions of silence in each audio file you open. Here you can set some settings to help AudioSlicer determine what is silence and what isn't.

• Minimum Silence Duration	Suppose you want to split words up in an audio file. In-between each word is silence, and each word itself is audio content. If you set the minimum silence duration too high, AudioSlicer might not notice the silence between two words and will consider them as one word. However, if you set the silence too low, AudioSlicer may start putting split points in the middle of words. For instance, the word airport could turn into "air <split point=""> port." The value you set here becomes the minimum value in the "Show Slices in Range" slider in the split document window.</split>
Minimum Silence Volume	AudioSlicer uses this value to distinguish be- tween ambient noise (silence) and audio con-

tent. If you set this value too low, AudioSlicer will consider pops, chair squeaks, and fan noise to be audio content. If you set it too high, you'll find the ends of your words getting cut off.

Export

• Relative Split Point in Silence

When you export all the slices as MP3 files, AudioSlicer needs to know where one slice stops and the other begins. This is called the split point. Normally the split point is in the middle of a silence. By adjusting the Relative Split Point in Silence value, you can change where one slice starts and the other begins. If you set the Relative Split Point in Silence to 50%, each split point will occur in the exact middle of a silence. If you set the Relative Split Point in Silence to 75%, each split point will occur three quarters of the way through the silence. Now if you export the slices, and play each file, the silence on the end of each file will be longer than the silence at the beginning.

Split Document Window

You'll do most of your work in the split document window. You can think of it as your project window. This is where you can adjust your audition settings and choose where you want to split your file.

Toolbar

The toolbar contains some handy commands that you might use often as well as a volume control.

• <i>i</i> Inspector Button	Opens and closes the inspector window.
• Forward and Back Buttons	The forward button moves to the next segment or track in the list while the back button moves to the previous one.
Volume Control	Adjusts the playback volume. This does not effect the exported files.
• Export Button	Exports each slice as a file.

Play Settings

Play setting do not effect your exported slices—they help you make the slices themselves.

• Play Interval Fields The play interval fields let you fine tune how much audio you want to listen to before and after a split point. You can setup defaults for these fields in the preferences window.

• Silence Range Slider

You can use the silence range slider to narrow down the list of silences you have to choose from. The first value is the shortest silence that will be displayed. So if it has a value of one second, any silence shorter than one second won't be displayed, and you won't be able to split there. The second value is the longest silence that will be displayed. If you have a 50 second silence and a maximum value of 3 seconds, some of that 50 second silence will be considered audio content.

Track and Slice List

The track and slice list shows all the tracks and segments in your split file. A segment can be either audio content or a silence. Silences have a length indicator and a split or join button. All segments including tracks have a play button.

At the top of the list is the heading row. In it are the labels for all the attributes of a track or segment. Some of the labels such as artist and genre only apply to tracks. These attributes can be edited in the inspector palette.

You can rearrange the labels by dragging them. You'll see a gray outline of the column move as you drag it. You can also change the column width by dragging the dividers between columns.

To the far right of the labels is a triangular shaped drop-down menu. From this menu you can check which columns you want displayed.

Inspector Window

The inspector window shows you information about whichever track or segment you have selected.

Overview

Overview	
• Selection	Shows how many slices or segments you have selected. Here slices are considered the same as tracks. If you select two tracks, the inspector window will show that you have two slices selected. If you select segments only then the inspector will report how many segments you have selected.
• Duration	Shows the total duration of all the slices or segments you have selected. Note that when you select a slice the duration will be the same no matter if you select its segments or not. This is because slices encompass seg- ments.

Details

The details section of the inspector palette is basically an ID3 tags editor for the slice or track you have selected.

Appendix A: Keyboard Shortcuts

Application Menu

Command	Shortcut
Preferences	¥,
Hide AudioSlicer	ЖН
Hide Others	жтн
Quit	жQ

File Menu

Command	Shortcut
Open	жo
Close	₩W
Save	₩S
Save As	¥企Q
Export Splitted	ЖЕ

Edit Menu

Command	Shortcut
Undo	₩Z
Redo	₩ûZ
Cut	ЖХ
Сору	ЖC
Paste	₩V
Select All	ЖA

Slices Menu

Command	Shortcut
Split Slices	ЖK
Join Slices	жJ
Play Previous Silence	¥-
Play Next Silence	¥+
Show Info	жı

Window Menu

Command	Shortcut
Minimize	₩M

Help Menu

Command	Shortcut
AudioSlicer Help	# ?

Legend: \mathfrak{K} : Command $\hat{\mathfrak{O}}$: Shift \mathfrak{T} : Option

∕⊠: Delete

17

Appendix B: FAQ

Q: Why can't I use system software prior to OS 10.3?

A: AudioSlice is based on Cocoa-Bindings which are only available in OS 10.3 and later. This is why Bernd won't be making a version for system software prior to OS 10.3. But hey, it's open source—maybe someone will try porting it to OS 10.2, but I seriously doubt it.

Q: Can I add more than one MP3 file to my split file?

A: No. The track listing in the split file window is for all your slices. You can think of tracks and slices as being roughly synonymous. Suppose you have several songs in a single MP3 file and you want to split them into tracks. Once you make the splits, you can assign ID3 tags to each track (or slice).

Appendix C: Troubleshooting