

Using Help

About Help

Adobe Systems Incorporated provides complete documentation in an Adobe PDF-based help system. This help system includes information on all tools, commands, and features of an application. It is designed for easy on-screen navigation and can also be printed and used as a desktop reference. Additionally, it supports third-party screen-reader applications that run in a Windows environment.

Navigating in Help

Help opens in an Adobe Acrobat window with the Bookmarks pane open. (If the Bookmarks pane is not open, click the Bookmarks tab at the left edge of the window.) At the top and bottom of each page is a navigation bar containing links to this page (Using Help), the table of contents (Contents), and the index (Index).

To move through pages sequentially, you can click the Next Page ▶ and the Previous Page ◀ arrows; click the navigation arrows at the bottom of the page; or click Back to return to the last page you viewed.

You can navigate Help topics by using bookmarks, the table of contents, the index, or the Search (Acrobat 6) or Find (Acrobat 5) command.

To find a topic using bookmarks:

- 1 In the Bookmarks pane, click the plus sign (+) (Windows) or the right-facing arrow (Mac OS) next to a bookmark topic to view its subtopics.
- 2 Click the bookmark to go to that topic.

To find a topic using the table of contents:

- 1 Click Contents in the navigation bar.
- 2 On the Contents page, click a topic to go to that topic.
- 3 To view a list of subtopics, click the plus sign (+) (Windows) or the right-facing arrow (Mac OS) next to the topic name in the Bookmarks pane.

To find a topic using the index:

- 1 Do one of the following:
 - Click Index in the navigation bar, and then click a letter at the top of the page.
 - In the Bookmarks pane, expand the Index bookmark to view the letter subtopics; then click a letter.
- 2 Locate the entry you want to view, and click the page number to go to that topic.
- 3 To view other entries for the same topic, click Back to return to the same place in the index, and then click another page number.



To find a topic using the Search command (Acrobat 6):

- 1 Choose Edit > Search.
- 2 Type a word or phrase in the text box and click Search. Acrobat searches the document and displays every occurrence of the word or phrase in the Results area of the Search PDF pane.

To find a topic using the Find command (Acrobat 5):

- 1 Choose Edit > Find.
- 2 Type a word or phrase in the text box and click Find. Acrobat searches the document, starting from the current page, and displays the first occurrence.
- 3 To find the next occurrence, choose Edit > Find Again.

Printing Help


Although Help is optimized for on-screen viewing, you can print selected pages or the entire file.

To print Help:

Choose File > Print, or click the Print icon in the Acrobat toolbar.

Contents

[Learning about Adobe Audition 4](#)
[What's New in Adobe Audition 1.5 7](#)
[Tutorials 9](#)
[Looking at the Work Area 19](#)
[Setting up Adobe Audition 37](#)
[Importing, Recording, and Playing Audio 56](#)
[Editing Audio 73](#)
[Enhancing and Restoring Audio 101](#)
[Applying Stereo, Pitch, and Delay Effects 138](#)
[Mixing Multitrack Sessions 179](#)
[Using Loops 214](#)
[Working with Video 222](#)
[Creating Surround Sound 226](#)
[Saving, Exporting, and Closing Files 237](#)
[Scripting and Batch Processing 249](#)
[Burning Audio CDs 259](#)
[Keyboard Shortcuts 264](#)
[Digital Audio Primer 267](#)
[Glossary 273](#)
[Legal Notices 284](#)



Learning about Adobe Audition

Getting help

There are a number of ways to get the help you need in Adobe Audition. The following three tables can help you find specific resources related to Adobe Audition features, training resources, and support.

Finding Help for Adobe Audition features	
If you ...	Try this ...
Want information about installing Adobe Audition	<ul style="list-style-type: none"> • Insert the Adobe Audition application CD into your CD drive, and follow the on-screen installation instructions. (You cannot run Adobe Audition from the CD.) • See the ReadMe file on the application CD.
Are new to Adobe Audition and want an overview of tools and features	<ul style="list-style-type: none"> • For information about specific tasks, see “Working with Adobe Audition” on page 6. • For information about the user interface, see “About the work area” on page 19. • Move the pointer over tools and buttons to view tool and button names. • See the beginning tutorials in Help.
Are upgrading from a previous version of Adobe Audition	See “What’s New in Adobe Audition 1.5” on page 7 to get an overview of new features. Or, for more detailed information, see the NewFeatures.pdf file on the Adobe Audition application CD.
Are looking for detailed information about a feature	<ul style="list-style-type: none"> • In Help, use the Index or Search tabs. • In windows and dialog boxes, click the Help button or press F1.
Want a list of keyboard shortcuts	See “Keyboard Shortcuts” on page 264 .



Finding Adobe Audition training resources	
If you ...	Try this ...
Want to obtain in-depth Adobe Audition training	<ul style="list-style-type: none"> • See the tutorials on the Adobe Studio Web site at www.studio.adobe.com. • Browse the Adobe Press materials at www.adobepress.com (English only) and the training resources at www.adobe.com/support/training.html. • For step-by-step lessons, consider the Adobe Classroom in a Book series.
Are looking for background information on digital audio	See the “Glossary” on page 273 and “Digital Audio Primer” on page 267 .
Want information about becoming an Adobe Certified Expert	Visit the Partnering with Adobe Web site at http://partners.adobe.com . Certification is available for several different geographical regions.
Want training from an Adobe Certified Training Provider	See the Training page of the Adobe Web site at www.adobe.com/support/training.html .

Finding support for Adobe Audition	
If you ...	Try this ...
Want customer or technical support	<ul style="list-style-type: none"> • Refer to the technical support card provided with your software. • See the Adobe Audition support page at www.adobe.com/support/products/audition.html. • See the ReadMe file installed with Adobe Audition for information that became available after this guide went to press.
Want answers to common troubleshooting questions	Visit the Adobe Audition support page at www.adobe.com/support/products/audition.html .
Want to register your copy of Adobe Audition	<ul style="list-style-type: none"> • When you first start Adobe Audition, you're prompted to register online. Fill out the form, and then submit it directly or fax a printed copy. • Fill out and return the registration card included with your software package.
Want to access downloads or links to user forums	Visit the main Adobe Audition page at www.adobe.com/audition .

Working with Adobe Audition

You can work with Adobe Audition in many different ways. In the related topics, you'll find directions to specific information to help you accomplish some common Adobe Audition tasks.

If you want to increase productivity

- Use the Organizer window to quickly organize files, preview loops, and apply effects. (See [“Organizing files and effects” on page 30](#) and [“Previewing audio by using the Organizer window” on page 69](#).)
- Automatically convert audio from a CD into an editable waveform. (See [“Importing audio from CD” on page 58](#).)
- Store selections and start points in cues to speed up editing and navigation tasks. (See [“Working with cues” on page 82](#).)
- Batch process files to quickly apply favorite processing or prepare audio for specific mediums, such as audio CD or the Web. (See [“Scripting and Batch Processing” on page 249](#).)

If you want to create video soundtracks

- Easily create and remix soundtracks used in Adobe® Premiere® Pro and After Effects® projects. (See [“Working with Adobe Premiere Pro and After Effects” on page 222](#).)
- Time stretch audio clips to match video. (See [“Time stretching audio clips” on page 195](#).)
- Generate noises and tones for sound effects. (See [“Generating audio” on page 89](#).)
- Create surround-sound mixes. (See [“About surround sound” on page 226](#).)

If you want to record and mix musical compositions

- Nondestructively record and edit multitrack sessions of up to 128 tracks. (See [“About mixing multitrack sessions” on page 179](#).)
- Automate mixes with clip envelopes. (See [“Automating mixes with clip envelopes” on page 207](#).)
- Apply, edit, and rearrange real-time effects, without making any permanent changes. (See [“Using real-time effects” on page 205](#).)
- Build compositions with musical loops. (See [“About loops” on page 214](#).)
- Synchronize with ReWire and SMPTE. (See [“Setting up ReWire connections” on page 42](#) and [“Setting up for SMPTE synchronization” on page 40](#).)

What's New in Adobe Audition 1.5

Use integrated tools

Adobe Audition tightly integrates with flexible audio technology like ReWire and VST, and video applications like Adobe Premiere Pro and Adobe After Effects.

ReWire support Stream full-resolution audio data in real-time between Adobe Audition and other audio software such as Propellerhead Reason and Ableton Live. (See [“Setting up ReWire connections” on page 42.](#))

VST plug-in support Expand your options with integrated support for third-party VST plug-ins, which can also be used in Adobe Premiere Pro. (See [“Using plug-in effects” on page 36.](#))

Enhanced video integration Edit video soundtracks with ease. View video frames in the track display, and import a wide range of video file formats, including AVI, MPEG, and WMV. (See [“About working with video” on page 222.](#))

Improved workflow with other Adobe products Work smoothly with Adobe Premiere Pro, Adobe After Effects, and Adobe® Encore™ DVD by using similar tools, menus, and keyboard shortcuts. (See [“Working with Adobe Premiere Pro and After Effects” on page 222.](#))

Sound your best

With high fidelity, 32-bit internal processing, Adobe Audition supports up to 32-bit files and sample rates up to 10 MHz. Powerful effects, restoration, and pitch correction tools let you create the exact sound you're after.

Pitch correction tool Correct off-pitch performances and create pitch-based effects. Use automatic mode for quick results, or manual mode for precise control. (See [“Using the Pitch Correction effect \(Edit View only\)” on page 152.](#))

Frequency space editing Visually isolate, select, and modify sounds in frequency and time using the Marquee Selection tool. (See [“Selecting audio frequencies in Spectral View” on page 77.](#))

Automatic elimination of clicks and pops Quickly and easily restore digital recordings of vinyl source material, wireless mics, DV cameras, and other production audio. (See [“Using the Auto Click/Pop Eliminator effect \(Edit View only\)” on page 108.](#))

Studio reverb Apply high-quality reverb that conserves processing resources, while offering extensive controls. (See [“Using the Studio Reverb effect” on page 168.](#))

New sample sessions Get up to speed quickly by using any of the 20 sample sessions included with Adobe Audition. Modify the samples to create your own music. (See [“About mixing multitrack sessions” on page 179.](#))



New royalty-free loops Use more than 500 new music loops—for a total of more than 5,000—in a variety of styles including 70's disco, classic rhumba, and wedding and event. (See [“About loops” on page 214.](#))

Vocal extraction Quickly and easily extract the vocal portions of a track to create either a cappella or karaoke-ready tracks, while preserving the stereo image. (See [“Using the Center Channel Extractor effect” on page 142.](#))

Flexible envelope scaling Rescale control points on pan, volume, and effects envelopes to quickly modify a clip in a multitrack mix. Scale all points simultaneously while maintaining relative or absolute relationships between points. (See [“Automating mixes with clip envelopes” on page 207.](#))

Work efficiently

Adobe Audition puts all the tools you need at your fingertips so you can get your work done quickly and efficiently. An intuitive interface gets you up and running in no time, and integrated editing, mixing, and CD burning streamline your audio workflow.

Integrated CD burning Create masters of your audio compositions by burning gapless audio CDs directly from Adobe Audition. (See [“Using CD Project View” on page 259.](#))

Time stretching Visually drag the edge of any audio clip in a multitrack mix to fit a specific length of time, with or without affecting the clip's pitch. Quickly fit sound effects and dialog to video clips. (See [“Time stretching audio clips” on page 195.](#))

Preroll and postroll playback Speed the process of performing destructive edits and applying effects by listening to the audio preceding and following a selection. (See [“Playing audio by using the transport controls” on page 67.](#))

Custom keyboard shortcut sets Customize keyboard shortcut sets to configure Adobe Audition for your working style. (See [“Using shortcuts” on page 21.](#))

In-time loop previews Use the Organizer window to preview loops in the tempo and pitch of the current session before adding them to your mix. (See [“Previewing audio by using the Organizer window” on page 69.](#))

Task-based documentation Quickly learn how to complete audio production tasks using an updated Help system and user guide organized by subjects such as editing, looping, and video.

Tutorials

About these tutorials

The following tutorials provide a quick tour of basic Adobe Audition 1.5 features. Before you get started, familiarize yourself with the basic concepts of Adobe Audition by looking at [“About the work area” on page 19](#).

As you work through these tutorials, you may have new ideas and questions. To assist you with learning Adobe Audition, Adobe provides a variety of resources. For more information, see [“Getting help” on page 4](#).



For more advanced Audition tutorials, visit the Adobe Web site at www.adobe.com/products/tips/audition.html.

Edit a voiceover and remove background noise

In Edit View, you can quickly import and edit audio files from a variety of sources, such as WAV files, audio CDs, or your own recordings. Powerful editing tools let you precisely remove unwanted sections of a performance, delete silence, and append multiple files. If an audio file contains noise, you can easily remove it by using several noise reduction effects.

In this tutorial, you'll use Edit View to combine two files and edit a voiceover. If you're not sure about the difference between Edit View and Multitrack View, take a minute to read [“About using Edit View and Multitrack View” on page 20](#) and [“Switching between views” on page 20](#).

1. Import an audio file

In Edit View, choose File > Open, and select the first file you want to use as a voiceover. If you're not sure which file you want, use the Play controls in the Open dialog box to preview files. Just click Play to hear the selected file; select Loop if you want the file to repeat until you click Stop, or click Auto Play to automatically play a file when you select it.


If you don't have a suitable file on hand, you can easily record one in Adobe Audition or extract one from a CD. For more information, see [“Recording audio” on page 64](#) and [“Extracting tracks from CDs” on page 59](#).



2. Remove unwanted audio

In the waveform display, drag to select the range of audio you want to remove. Then choose Edit > Delete Selection. Or, select the segment of audio you want to keep, and choose Edit > Trim.

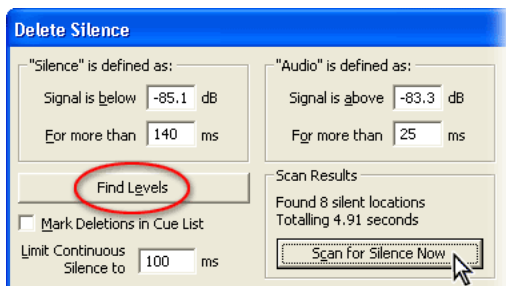
In this example, we used the Trim command to remove unwanted audio from both sides of the selected audio.



3. Delete silence

Choose Edit > Delete Silence. To automatically locate a good starting point for the Silence and Audio signal levels, click Find Levels. Click Scan For Silence Now to preview the amount of silence that will be removed, as well as the number of silence sections that were found. If you're satisfied with the results, click OK. Otherwise, adjust the Delete Silence options and click Scan For Silence Now again.

Deleting silence is useful when you want to clean up voice prompts and speed up narratives without affecting the foreground audio. In addition to deleting silence, you can also generate it to insert pauses into an audio file. For example, you may want to add silence to the end of a waveform before appending another file. For more information, see ["Deleting silence" on page 88](#).



4. Append another audio file

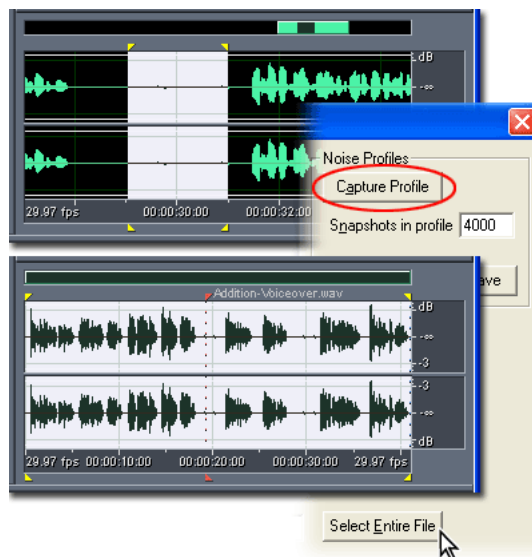
You can easily combine two audio files by choosing File > Open Append. Select the file you want to append to the end of the current file, and click Append.



5. Remove background noise

If your voiceover contains background noise, you'll hear it during moments when there should be silence. To remove this noise, select a range of audio that contains only noise and is between one-half second and two seconds long (a two-second sample works best). Then choose Effects > Noise Reduction > Noise Reduction, and click Capture Profile. Adobe Audition automatically determines the best settings for removing the noise based on this profile.

Next, click Select Entire File to apply the noise reduction to the entire waveform. Click Preview to listen to the results, and click Bypass to compare the preview to the original audio. If you're satisfied with the results, click OK. If you're not satisfied with the results, experiment with the noise reduction settings. For more information on specific noise reduction options, see ["Noise Reduction options" on page 117](#).



6. Save the file

When you're satisfied with the audio, choose File > Save As. Specify a location for the file, type a filename, and choose a file format. Depending on the format you choose, additional options may be available. To view and specify format-specific options, click Options. When you are finished setting options, click Save.

In general, Windows PCM format is the best choice for saving uncompressed audio, and either mp3PRO® or Windows Media Audio is the best choice for saving compressed audio. For more information on the different audio formats supported by Adobe Audition, as well as format-specific options, see ["Choosing an audio file format" on page 240](#).

Mix the voiceover with background music


In Multitrack View, you can mix together multiple audio files to create layered soundtracks and elaborate musical compositions. Because mixing occurs in real time and is nondestructive, it's extremely flexible; during playback, you can adjust mixes and record additional tracks without making any permanent changes. If a mix doesn't sound good next week, or even next year, you can simply remix the original audio files.

In this tutorial, you'll use the edited voiceover from the previous tutorial as a starting point. Then you'll record background music on a separate track, apply a real-time effect, and mix down to a new audio file.

1. Insert an existing audio file into a track

Above the waveform display, click the Multitrack View tab. In the track controls for track 1, type **Voiceover** in the name text box. In the track display, position the current-time indicator at the beginning of the multitrack session. Then choose Insert > Audio, and select the voiceover file you edited in the previous tutorial. The file becomes an audio clip in the selected track.

If an audio file is currently open in Edit View, you don't need to use the Insert > Audio command. Instead, simply choose the file from the recent items list at the bottom of the Insert menu.






2. Configure another track for recording

In the track controls for track 2, type **Music** in the name text box, click the Record-enable button **R**, and then click the In button. In the Input Devices dialog box, select a sound card input that you've attached a microphone to, and click OK.

If your sound card has multiple inputs, you can use this process to simultaneously record on multiple tracks. For example, you can record the performance of a live band and place drums, keyboards, and guitars on separate tracks. Alternatively, if your sound card has few inputs or a group performance is impractical, you can sequentially record each instrument on a separate track.



3. Record the background music

To prevent microphone feedback, turn the volume of your monitor speakers completely down, and put on headphones. In the Transport Controls window, click the Record button , and play the background music while listening to the voiceover. When you're finished, click the Stop button  in the Transport Controls window. The new recording becomes an audio clip in the selected track. (To avoid inadvertently recording over the clip, again click the Record-enable button  in the track controls; the clip loses its red color.)

If you don't like a section of your musical performance, you can record over it by using the Punch In command. (See ["Recording audio in Multitrack View" on page 66.](#))




4. Apply a real-time effect

In the session display, select the Music track. In the Effects tab of the Organizer window, expand Delay Effects, and double-click Studio Reverb. In the Studio Reverb tab of the FX dialog box, choose Great Hall from the preset pop-up menu. Then adjust the balance of original and reverberant sound by dragging the Original and Reverb sliders.

When choosing a reverb effect, you must make trade-offs between sound quality and processing load. For example, the Full Reverb effect has a very rich sound, but it requires significant processing, reducing the performance of Adobe Audition. By contrast, the Studio Reverb effect works well for most reverb tasks and requires very little processing.



5. Edit the mix

In the Transport Controls window, press the Loop button . As the session plays, adjust the track volume and pan in the track controls, and refine the balance of original and reverberant sound in the FX dialog box. To begin the voiceover after the background music, click the Move/Copy tool in the toolbar, and drag the voiceover clip to the right.

Mix settings don't change audio files, so you can freely experiment with new settings until you arrive at the right sound. You can even automate mix settings over time by using clip envelopes. (See ["Automating mixes with clip envelopes" on page 207.](#))



6. Mix down to a new audio file

When you're happy with the sound of the mix, save the session, and then export an audio mixdown file. To save the session, choose File > Save Session, enter **FirstSession.ses** for the filename, and click Save. To export an audio mixdown, choose File > Export > Audio, specify a filename and format, and click Save.

From one session file, you can create multiple audio mixdowns, each optimized for a different medium, such as radio, video, or the Web. You can export audio mixdowns in a variety of common formats, ranging from WAV to mp3. Later, you can edit exported mixdown files in Edit View or burn them onto audio CD's. You can even incorporate them into other sessions, using the real-time controls in Multitrack View to quickly build layered musical compositions and video soundtracks.

Enhance a video soundtrack

Adobe Audition tightly integrates with digital video, enhancing any video project with professional sound. To create uniquely sophisticated soundtracks, you can combine Adobe Audition with Adobe Premiere Pro and Adobe After Effects to take full advantage of Adobe Audition's flexible mixing features.

In this tutorial, you'll create a separate mixdown file that is optimized for a video. To do so, you'll add a video clip to the multitrack session you created in the previous tutorial. Then you'll enhance and remix a video soundtrack from an Adobe Premiere Pro or After Effects project.

1. Insert a video file into the session

Open FirstSession.ses, the session you saved in the previous tutorial. In Multitrack View, select track 3 and type **Video** in the track's name text box. Then, type **Ambient** in the name text box for track 4. Select track 3 and drag the current-time indicator to the point at which you want to insert the video file.

Now, choose Insert > Video and select an AVI video file exported from Adobe Premiere Pro or Adobe After Effects. The video file is inserted into the selected track you named "Video," and its accompanying audio is inserted into the track below it named "Ambient." In our example, the audio track contains ambient nature sounds recorded when the video was made. (If the audio you're inserting was saved at a different sample rate from your session, Adobe Audition notifies you and lets you convert it to the correct rate.)

Notice that a separate Video window opens in Adobe Audition, containing the images from the video file. This window lets you view the video as you play the audio, helping you mix and synchronize tracks. Don't worry, though—the original video remains unchanged in your Adobe Premiere Pro or After Effects project.



2. Reposition and time stretch audio clips to match video

Drag each of the clips in the tracks from the previous session so that they align with the start of the Video track and Ambient track you just inserted. Now all the tracks will begin at the same time. If the audio clips from the previous session end too early (as they do in our example), stretch them out so that they synch up with the end of the video clip.

Hold down the Ctrl key and drag the lower right corner of each audio clip so that it extends to the end of the video clip. Time stretching lets you lengthen the clip without changing the pitch of the sounds within it. (See [“Time stretching audio clips” on page 195.](#))



3. Preview and mix the enhanced soundtrack

Click the Play button ▶ to preview the soundtrack. If the audio doesn't synchronize to the video or if it doesn't sound the way you want, edit it, adjust the audio tracks, and preview it again until you're satisfied with the result. When you're finished, save the session.

For instance, you can increase the volume of the Voiceover track by dragging to the right in the V text box. You can also change the reverb on the Music track by clicking the track's FX button to display the applied Studio Reverb effect. Or you can time stretch the audio clip for the Voice track to make it end a bit earlier.



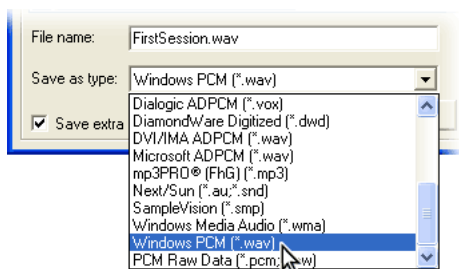
To change mix settings during video events such as scene changes, use automated clip envelopes. (See [“Automating mixes with clip envelopes” on page 207.](#))



4. Export the enhanced soundtrack

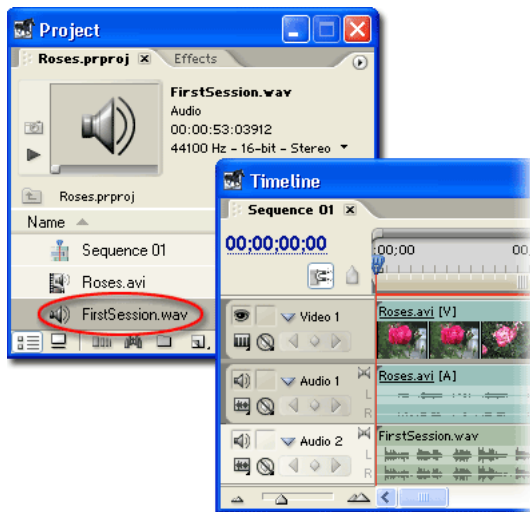
You're now ready to export the audio to your video application. Because you want to export the entire audio session, make sure that nothing is selected in the track display. (If necessary, click the track display to reveal the current-time indicator.) Choose Options > Settings, click the Data tab, and select Embed Project Link Data For Edit Original Functionality. Then choose File > Export > Audio. Name the file, choose WAV format for the file type, and click OK.

Adobe Audition lets you export either audio or video mixdowns. If you are producing final video output using a video-editing application (such as Adobe Premiere Pro) you'll create an audio mixdown. If you don't have a video-editing application, however, you can export video mixdowns that combine stereo audio with video. In this tutorial, we chose the WAV format because it lets us easily re-edit the file in Adobe Audition from Adobe Premiere Pro or After Effects.



5. Import the enhanced soundtrack in Adobe Premiere Pro or After Effects

In Adobe Premiere Pro or After Effects, open your video project. Then choose File > Import, and import the audio mixdown file you created in Adobe Audition.




6. Remix the video soundtrack in Adobe Audition

Play the video and audio in Adobe Premiere Pro or After Effects, and then edit your project as desired. If you need to change the soundtrack to reflect your video edits, you can remix it easily in Adobe Audition. To do so, simply select the mixdown file in the Adobe Premiere Pro or After Effects project, and choose **Edit > Edit Original**. Select **Launch The Audition Multitrack Session Which Created This File**, and click **OK**.

The session opens in Adobe Audition, where you can make any sound edits you want. For example, to make the Voiceover track stand out more, increase the midrange by clicking the **EQ** tab in the track controls and dragging to the right in the **M** text box. Then lower the volume of the Music track.


When you're done making changes, resave the session, and then reexport the audio mixdown file. Adobe Premiere Pro or After Effects automatically detects the newly saved file and updates it in the video project.



Looking at the Work Area

About the work area

Adobe Audition is divided into three main work areas: Edit View, Multitrack View, and CD Project View. This division is intended to help you focus on the major tasks of editing audio files, mixing sessions, and burning CDs. For more information on the differences between Edit View and Multitrack View, see [“About using Edit View and Multitrack View” on page 20](#). For more information on CD Project View, see [“Using CD Project View” on page 259](#).



Adobe Audition work area

A. Edit View tab **B.** Multitrack View tab **C.** CD Project View tab **D.** menus **E.** toolbars
F. display window **G.** various windows

All three views have a similar user interface, including the following components:

Menus The menus in the menu bar contain commands for performing tasks. (See [“Choosing commands” on page 21](#).)

Toolbars The toolbars hold buttons for applying commonly used functions. (See [“Using toolbars” on page 22](#).)

Windows Windows—including the Organizer, Transport Controls, Zoom Controls, Level Meters, and Selection/View Controls—help you monitor and modify audio files. (See [“Using windows” on page 22](#).)



Display window The display window shows you sound in an easy-to-manipulate form. In Edit View, the display window is where you modify single waveforms. In Multitrack View, the display window is where you mix multiple audio files in a session. (See [“About editing audio” on page 73](#) and [“About mixing multitrack sessions” on page 179](#).)



You can change many aspects of Adobe Audition’s appearance, including the color scheme, the appearance of buttons, and the appearance of the waveform display, in the Settings dialog box. (See [“Setting Adobe Audition preferences” on page 43](#).)

About using Edit View and Multitrack View


Adobe Audition provides different work areas for editing single waveforms and creating multitrack mixes. To edit single waveforms, you use Edit View. To mix multiple waveforms with MIDI and video files, you use Multitrack View.

Edit View and Multitrack View use different editing methods, and each has unique advantages. Edit View uses a *destructive* method, which changes audio data, permanently altering saved files. Such permanent changes are preferable when converting sample rate and bit depth, mastering, or batch processing. Multitrack View uses a *nondestructive* method, which is impermanent and instantaneous, requiring more processing power, but increasing flexibility. This flexibility is preferable when gradually building and reevaluating a multilayered musical composition or video soundtrack.

You can combine destructive and nondestructive editing to suit the needs of a project. If a multitrack clip requires destructive editing, for example, simply double-click it to access Edit View. Likewise, if an edited waveform contains recent changes that you dislike, use the Undo command to revert to previous states—destructive edits aren’t applied until you save a file. For more information on using Edit View, see [“About editing audio” on page 73](#); for more information on using Multitrack View, see [“About mixing multitrack sessions” on page 179](#).

Switching between views

You can use the tabs above the display window or menu commands to switch between Edit View, Multitrack View, and CD Project View. If you prefer not to use the tabs above the display window, you can hide them.







View tabs above the display window

To switch between views:

Do one of the following:

- Choose View > Edit Waveform View, View > Multitrack View, or View > CD Project View.
- Click the Edit View tab, the Multitrack View tab, or the CD Project View tab above the display window.

- Click the Edit Waveform View button , Multitrack View button , or CD Project View button  in the View toolbar. (See [“Using toolbars” on page 22.](#))
- In Multitrack View, double-click a file in the Files tab of the Organizer window or select a file and click the Edit File button . Alternatively, double-click a waveform block in the display window.

To show and hide view tabs above the display window:

Choose View > Show View Tabs. A check mark indicates that the tabs are showing.

Choosing commands


Commands let you perform a wide variety of tasks. You can choose commands from the menus at the top of your screen or click buttons in a toolbar. You can also use context-sensitive (right-click) menus and keyboard shortcuts to quickly execute commands.

Using context-sensitive menus

Adobe Audition makes liberal use of context-sensitive menus. Whenever you see a simple function button, control, window, or waveform action, try right-clicking it. Chances are you'll be surprised by a useful shortcut menu or a set of handy options that can make Adobe Audition's operation even easier.

Using shortcuts

Adobe Audition provides a set of standard keyboard shortcuts to help you speed up the editing process. For example, instead of using your mouse to go to the Edit menu and choose the Cut command, you can simply press Ctrl + X to cut the selected portion of a waveform. When available, the keyboard shortcut appears to the right of the command name in the menu or in the tool tip for a button or icon. Adobe Audition also provides keyboard shortcuts for performing certain mouse actions. These shortcuts are listed in the Keyboard Shortcuts appendix.

 If a shortcut isn't working, it's likely that the window you're trying to run the shortcut in doesn't have focus. For example, if you're in Edit View and you push F11 to bring up the Convert Sample Type dialog box and nothing happens, the waveform display probably isn't the active window. Click the waveform display to give it focus, and then try the shortcut again.

You can change nearly all of the default shortcuts and add shortcuts for other functions. In addition, you can add shortcuts that let you execute commands using keys on a MIDI keyboard, a sequencer, or any other device capable of issuing a MIDI command. This type of shortcut is referred to as a MIDI Trigger. For example, you can assign the Play command in Adobe Audition to the C4 note on your MIDI keyboard.

To enable MIDI triggering:

Choose Options > MIDI Trigger Enable. A check mark indicates the MIDI triggering is on.

Important: Before attempting to enable MIDI triggering, you must choose a device for MIDI In that's recognized by Windows. For more information, see [“Designating which devices you want to use” on page 37.](#)

To customize a shortcut:

- 1 Choose Options > Keyboard Shortcuts And MIDI Trigger.
- 2 Select the function you want to assign the shortcut to.

Note: You can filter the list of functions by choosing an option from the Category menu and clicking the Multitrack View or Edit View button. To show all functions, choose (show all) from the Category menu, and deselect the Multitrack View and Edit View buttons.

- 3 Do any of the following:
 - To assign a keyboard shortcut to the function, click in the Keyboard Shortcut text box and press the desired keyboard combination. Many Adobe Audition users find single key shortcuts (such as *n* for Normalize) faster to use and easier to remember.
 - To assign a MIDI trigger to the function, click in the MIDI Trigger text box and press the desired key on the MIDI keyboard. You can also apply MIDI events other than pressing keys (such as pressing the foot pedal).
 - To remove a keyboard shortcut or MIDI trigger from the function, click Clear.
- 4 If you enter a key combination that's already in use, Adobe Audition notifies you of the conflict in the Conflicting Keys text box. Click Clear, and enter a different shortcut before continuing.
- 5 Click OK.

To restore the default keyboard shortcuts:

- 1 Choose Options > Keyboard Shortcuts And MIDI Trigger.
- 2 Choose Adobe Audition Default from the Set list, and click OK.

Using toolbars

Many of Adobe Audition's most commonly used functions are represented as buttons within toolbars, which appear near the top of the main interface. These buttons give you instant access to effects, file handling functions, viewing options, and more, at the press of a button.



To see what a button does, hold your mouse pointer over it to display a tool tip that describes the function in simple terms.

To show or hide a toolbar:

Choose View > Toolbars, and choose a toolbar name from the submenu. A check mark indicates that the toolbar is showing.

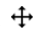
To specify how many rows of buttons are displayed:

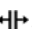
Choose View > Toolbars, and choose a number of rows from the submenu.


Using windows

Many windows in the Adobe Audition interface can be repositioned and resized to better suit your requirements. You can also hide windows that you're not currently using, and then show them again when needed. For more information on specific windows, see the index or search Help.

Repositioning and resizing windows

When you reposition a window, you can dock it in a specific location in the interface, or you can undock the window so that it floats above the main window. To identify docked windows, look for two thin vertical or horizontal lines. These lines are the *handle* (or grab bar) of a docked window. Move your mouse over a handle, and the cursor looks like a plus sign with arrows at each end .

Some docked windows can also be resized. If resizing is possible, the docked window will have a single, thicker horizontal or vertical bar, called a *resize bar*. When you move your mouse over a resize bar, the cursor takes on the appearance of two lines with two arrows .



Docked window

A. Handle **B.** Resize bar

To undock a window:

Drag the window's handle to the middle of the work area until you see an outline of the window.

The window is now a standard floating window. You can move the window by dragging its title bar.



Press Ctrl while moving a floating window around to force it to not dock. That way you can float the window over an area that it would normally try to dock to. To disable this feature, select Ctrl Key Allows Dockable Windows to Dock in the General tab of the Settings dialog box. (See ["Setting Adobe Audition preferences" on page 43.](#))

To dock a window in a different location:

1 Drag the window's handle around the work area to locate potential docking areas. The resize bars of other docked windows will light up wherever docking is possible.

2 When you locate the desired docking area, release the mouse button. The window snaps into its new location.



If a window is docked in the same row with other windows, you can force the window into a new row by right-clicking the window's handle and selecting Force New Row. Likewise, deselecting Force New Row causes the window to dock in the previous row (if there's room).

To resize a docked window:

Drag the window's resize bar.

Even if the resize bar is visible, resizing might not be possible due to the other windows that are in the row with the window you're trying to resize.



To reset windows to the default layout, select Restore Default Workspace in the General tab of the Settings dialog box. (See ["Setting Adobe Audition preferences" on page 43.](#))

Showing and hiding windows

You can free up space in the work area by closing windows when you aren't using them, and then redisplay the windows as needed. The Window menu lists all available windows; a check mark indicates that a window is currently showing.

To hide a window:

Do one of the following:

- Choose the window name from the Window menu.
- Click the button that corresponds to the window name in the View toolbar. (See [“Using toolbars” on page 22.](#))
- For docked windows, right-click the window's handle and choose Close.
- For undocked windows, click the X button on the window's title bar.

To show a window:

Choose the window name from the Window menu, or click the window's button in the View toolbar.

Using placekeeper windows

Placekeeper windows let you define the aspect ratio of a docking area. For example, if you try docking the Track EQ controls above the transport controls, they end up going underneath the whole session display, which creates a view that isn't very useful (or aesthetically pleasing). You can use a placekeeper, though, on either side of the Track EQ to force the EQ into a certain aspect ratio. You can also use placekeepers just for appearance's sake, just because you like the way they let you customize the work area.

You can create up to four placekeeper windows, and insert them wherever docking is allowed. You can also change the appearance of placekeeper windows by filling them with a pattern.

To insert a placekeeper window:

- 1 Choose Window > Placekeeper.
- 2 Dock the placekeeper in the desired location. The window is automatically resized to fit the docked area.

To change the appearance of a placekeeper window:

Right-click the window's handle, and choose a fill option: Nothing, Cool Texture, or Squares. To make future placekeeper windows adopt the current appearance, choose Make Default.

To delete a placekeeper window:

Right-click the window's handle, and choose Close.

Navigating in the display window

The display window shows you the current waveform (in Edit View) or session (in Multi-track View). You can control how much of the waveform or session is displayed by zooming and scrolling. You can also use the selection and view controls to determine the beginning time, ending time, and length of audio data in the display window.

Zooming

Zooming lets you adjust the view in the display window to best meet your needs. For example, you can zoom in to clearly see the samples in a waveform, or you can zoom out to get a visual overview of a waveform or session.


The Zoom Controls window provides a variety of tools for zooming. You can also zoom by dragging in the horizontal scroll bar, vertical scroll bar (Multitrack View only), or vertical ruler.



Zoom controls

To show or hide the zoom controls:





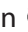



Do one of the following:

- Choose Window > Zoom Controls. A check mark indicates that the controls are visible.
- Click the Hide/Show Zoom Controls button  in the View toolbar. (See [“Using toolbars” on page 22.](#))

If you don't like the default location of the zoom controls, you can reposition them or detach them so they float above the main window. (See [“Using windows” on page 22.](#))

To zoom in or out by using the zoom controls:

Do any of the following:

- Click the Zoom In Horizontally button  to zoom in on the center of the visible waveform window or session.
- Click the Zoom In Vertically button  to increase the vertical scale resolution of a waveform's amplitude display (in Edit View) or decrease the number of viewed tracks in the session display (in Multitrack View).
- Click the Zoom To Selection button  to zoom in on the actively selected waveform or session range.
- Click the Zoom In To Right Edge Of Selection button  to zoom in on the right boundary of the actively selected waveform range or session.
- Click the Zoom In To Left Edge Of Selection button  to zoom in on the left boundary of the actively selected waveform range or session.
- Click the Zoom Out Horizontally button  to zoom out from the center of the visible waveform window or session.
- Click the Zoom Out Full Both Axis button  to zoom out to display the entire waveform or blocks that are contained within a session.
- Click the Zoom Out Vertically button  to decrease the vertical scale resolution of a waveform's amplitude display (in Edit View) or to show more tracks in the session display (in Multitrack View).

To zoom in or out by using a scroll bar or ruler:

Do either of the following:


- To change the viewable range of time, position the pointer in the timeline or over the left or right edge of the horizontal scroll bar. Then drag to the left or right. A magnifying glass with arrows icon appears as you drag.
- To change the viewable range of amplitude (in Edit View) or tracks (in Multitrack View), hold down the right mouse button in the vertical ruler, and drag up or down. The magnifying glass with arrows icon appears as you drag.



You can also use the wheel on your mouse to zoom in and out. To do so, place the pointer over the horizontal scroll bar, timeline, vertical scroll bar (Multitrack View only), or vertical ruler, and roll the mouse wheel. To set a zoom percentage for the mouse wheel, enter a value for Zoom Factor in the General tab of the Settings dialog box. (See [“Setting Adobe Audition preferences” on page 43.](#))

Scrolling

The display window provides several scrolling devices. The horizontal scroll bar—which, by default, is at the top of the display window—lets you scroll forwards and backwards in time throughout a waveform (in Edit View) or session (in Multitrack View). The vertical ruler on the right side of the display window lets you scroll through amplitude ranges (in Edit View) or tracks (in Multitrack View). In Multitrack View, there’s an additional vertical scroll bar on the left side of the display window that lets you scroll through tracks.



Scrolling devices

A. Vertical scroll bar **B.** Horizontal scroll bar **C.** Vertical ruler

To scroll in the display window:

Do either of the following:

- To scroll to the left or right, drag the horizontal scroll bar. Or, click to the left or right of the scroll bar to page through the display one screen at a time.
- To scroll up or down, drag in the vertical ruler. In Multitrack View, you can also drag the vertical scroll bar or click above or below the scroll bar to page through the display one screen at a time.



You can also use the wheel on your mouse to scroll in the display window. To do so, place the pointer over the display window, and roll the mouse wheel.


To change the position of the horizontal scroll bar:

Right-click the horizontal scroll bar, and choose a display option: Above Display or Below Display.

Using the selection and view controls

The Selection/View Controls window shows the beginning and ending points, as well as the total length of both the selection and the section of the waveform or session that's currently visible. Both the selection and display range is shown in the current time-display format. For information on changing the time-display format, see [“Monitoring time” on page 63](#).


In addition to viewing time information, you can also use the selection and view controls to adjust selections and change the section of audio data that is visible in the display window. Simply enter new values for Begin, End, and Length. After you click in a text box, you can right-click to access additional context-menu commands.



Selection and view controls

To display the selection and view controls:


Do one of the following:

- Choose Window > Selection/View Controls. A check mark indicates that the window is showing.
- Click the Hide/Show Selection/View Controls button  in the View toolbar. (See [“Using toolbars” on page 22](#).)

If you don't like the default location of the selection and view controls, you can reposition them or detach them so they float above the main window. (See [“Using windows” on page 22](#).)

Using the status bar

The status bar runs along the very bottom of Adobe Audition's main window. It can display information such as sample format, file size, and free disk space.




Status bar

A. Data Under Cursor **B.** Sample Format **C.** File Size **D.** File Size (time) **E.** Free Space
F. Free Space (time) **G.** Keyboard Modifiers **H.** SMPTE Slave Stability

To show or hide the status bar:

Do one of the following:

- Choose View > Status Bar > Show. A check mark indicates that the status bar is visible.
- Click the Hide/Show Status Bar button  in the View toolbar. (See [“Using toolbars” on page 22.](#))

To change the type of information that is displayed in the Status Bar:

Choose View > Status Bar or right-click the Status Bar, and select the desired display options. Selected items appear in the Status Bar; unselected items are hidden.

You can choose from the following options:

Data Under Cursor Shows useful information such as the channel (if a current waveform is stereo), the amplitude (measured in decibels), and the time (hours:minutes:seconds:hundredths of seconds) from the beginning of the audio file. This data is computed at the precise point where your mouse pointer is placed within the wave display, and changes dynamically when you move the pointer. For example, if you see *R: -15.2 dB @ 0:00:242* in the Status Bar when in Edit View, this means that your pointer is over the right channel at 0.242 seconds into the waveform, and the amplitude at that precise point is -15.2 dB.

In the Multitrack View, you'll see even more beneficial data such as Pan and Volume envelope positions, envelope positions for effects envelopes, dynamic effect settings, and the current position of the wave block as you drag it around.

Sample Format Displays sample information about the currently opened waveform. For example, a 44,100 kHz 16-bit stereo file shows up as *44100 – 16-bit – stereo*.

File Size) Represents how large the active audio file is, measured in kilobytes. If you see *308 K* in the Status Bar, then the current waveform or session is 308 kilobytes (KB) in size.

File Size (time) Shows you the length (measured in time) of the current waveform or session. For example, *0:01:247* means the waveform or session is 1.247 seconds long.

Free Space In Edit View and Multitrack View, shows how much space is available on your hard drive. In CD Project View, shows how much space remains on a CD based on which View menu item is selected: 74 min CD or 80 min CD.

Free Space (time) In Edit View and Multitrack View, displays the amount of available time left for recording, based upon the currently selected sample rate. This value is shown as minutes, seconds, and thousandths of seconds. For example, if Adobe Audition is set to record an 8-bit mono waveform at 11,025 kHz, the time left might read something like *4399:15.527 free*. Change the recording options to 16-bit stereo at 44,100 kHz, and the remaining time value becomes *680:44.736 free*.

In CD Project View, shows how much space remains on a CD based on which View menu item is selected: 74 min CD or 80 min CD.

Keyboard Modifiers Displays the status of your keyboard's Ctrl, Shift, and Alt keys.


SMPTE Slave Stability Indicates the stability of incoming SMPTE timecode compared to Adobe Audition's internal clock. For example, *95.0% SMPTE* indicate a very strong SMPTE signal. Percentages above 80% should be stable enough to maintain sync. For more information on SMPTE synchronization, see [“Setting up for SMPTE synchronization” on page 40](#) and [“Using sessions as SMPTE masters or slaves” on page 183](#).

Undoing and redoing changes

Adobe Audition keeps track of the edits you perform during the course of an editing session. These changes are stored in a temporary file on your hard drive. They aren't permanently applied to the file until you save and close it, giving you unlimited undo and redo capability.

When you work with very large audio files, you might not have enough free disk space to save the Undo data before continuing with an edit. In addition, the time required to save the Undo information might slow down your work. You can solve either problem by disabling the Undo function.

To undo a change:

Choose Edit > Undo *[name of change]*. Or, click the Undo button  in the toolbar.

The Undo command conveniently indicates which change you're undoing. For example, it may appear as Undo Delete or Undo Normalize. If you haven't yet edited a waveform, or if Undo is disabled, this command appears as Can't Undo.




If you forgot which editing action you last performed on a waveform, look at the Undo command to refresh your memory, whether you want to undo the action or not.

To discard edits made since you last saved the file:

In Edit View, choose File > Revert To Saved.

To redo a change:

In Edit View, choose Edit > Redo *[name of change]*. Or, click the Redo button  in the toolbar.

To repeat the last command:

In Edit View, choose Edit > Repeat Last Command. You can repeat most editing functions in Adobe Audition by using this command; however, there are a few exceptions (such as Delete).

To disable or enable the Undo function:

Do one of the following:

- In Edit View, choose Edit > Enable Undo/Redo. A check mark indicates that the Undo function is enabled.
- Choose Options > Settings, and click the System tab. Select or deselect Enable Undo, and click OK. You can also specify the minimum number of undo levels, and you can purge all undo files. (See ["System options" on page 45.](#))



If you don't have enough disk space to save the undo information, you can change the Temp folder to a different drive, if available.

Organizing files and effects

The Organizer window appears in Edit View, Multitrack View, and CD Project View. This handy, tabbed window lets you easily open and close files, see a list of all open waveforms and MIDI files, choose effects with ease, and more. By default, the Organizer window is docked to the left of the waveform or session display; however, you can reposition it or detach it so it floats above the main window. (See [“Using windows” on page 22.](#))

Organizing files

The Files tab in the Organizer window displays a list of open waveforms, MIDI files, and video files. You can use the Files tab to import files, select files for editing, insert clips into sessions, insert tracks into CDs, and close files.

The Files tab also provides a variety of advanced options that let you show and hide cues, change the listing and sort order of files, and play files. You can choose to hide advanced options if you don't use them.



Files tab in the Organizer window

To display the Files tab:

- 1 If the Organizer window isn't showing, choose Window > Organizer to display it.
- 2 Click the Files tab in the Organizer window. The following buttons appear at the top of the Files tab:
 - The Import File button lets you import audio, MIDI, and video files into Adobe Audition.
 - The Close Files button lets you close all selected files in the Files tab.
 - The Insert Into Multitrack button lets you insert all selected files, each into their own track, in Multitrack View. (See [“Inserting audio files into multitrack sessions” on page 58.](#))
 - The Insert Into CD Project button lets you insert all selected files into CD Project View. (See [“Inserting tracks” on page 259.](#))
 - The Edit File button lets you open the selected file in Edit View. (See [“Switching between views” on page 20.](#))


To select files in the Files tab:

Do any of the following:

- To select a single file, click it.
- To select adjacent (or *contiguous*) files, click the first file in the desired range, and then Shift-click the last.
- To select nonadjacent (or *noncontiguous*) files, Ctrl-click them.

Note: If you select multiple files, only the last file you click appears in Edit View.

To show or hide advanced options in the Files tab:

Click the Advanced Options button  at the top of the Files tab. When showing, the advanced options appear at the bottom of the Files tab.

For information on the play controls in the Files tab, see [“Previewing audio by using the Organizer window” on page 69](#).

To change the listing and sort order of files in the Files tab:

Make sure that the advanced options are showing, and do any of the following:

- To show or hide files, select a Show File Types option. An X indicates that files of the specified type are showing.
- To change the sort order of files, choose an option from the Sort By menu.
- To display the full path *[drive, folder(s), filename]* of the entries in the File tab, select the Full Path button. To display only the filenames, deselect this button.

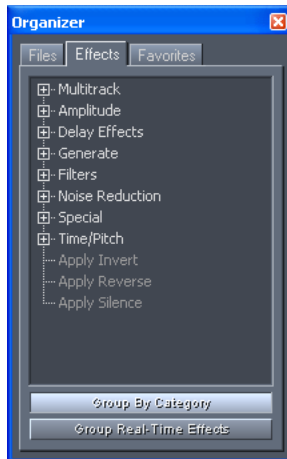
To show or hide cues in the Files tab:

Make sure that the advanced options are showing, and click Show Cues.

When Show Cues is selected, a plus icon appears next to files that contain cues. Click the plus icon to display the cue names. For more information on cues, see [“Working with cues” on page 82](#).

Organizing effects

The Effects tab in the Organizer window lists all of the effects at your disposal. The listing includes all of Adobe Audition's effects as well as all installed DirectX and VST audio plug-ins. You can change the grouping of effects to best meet your needs.



Effects tab in the Organizer window

To display the Effects tab:

- 1 If the Organizer window isn't showing, choose Window > Organizer to display it.
- 2 Click the Effects tab in the Organizer window.

To change how the effects are grouped:

Click the buttons at the bottom of the Effects tab:

- Select Group By Category to list effects in a hierarchy where categories and their entries are shown in the same order as they appear in the Effects menu.
- Deselect Group By Category to display all effects in roughly the same order as they appear in the Effects and Generate menus.
- Select Group Real-Time Effects to list effects in a hierarchy where all of the Real-Time Effects are grouped together, the Off-Line Effects are grouped together, and the Multi-track Effects are grouped together.
- Deselect Group Real-Time Effects to return to the previous view.

Organizing favorites

Favorites are effects, scripts, and even third-party tools that you've saved for easy access. The Favorites tab in the Organizer window lists all of the favorites you've created. (These same items are listed in the Favorites menu.)



Favorites tab in the Organizer window

To display the Favorites tab:

- 1 If the Organizer window isn't showing, choose Window > Organizer to display it.
- 2 Click the Favorites tab in the Organizer window.

For more information on creating and editing favorites, see [“Using favorites \(Edit View only\)” on page 256](#).

Working with effects

Effects provide much of the functionality in Adobe Audition. For example, you use effects to remove noise, optimize volume, change pitch, and add reverb. If Adobe Audition doesn't provide the effect you want, you may be able to purchase a plug-in effect to do the job.

As you apply effects, you'll notice similarities between Adobe Audition's effect dialog boxes. For example, many effect dialog boxes provide presets for storing and recalling your favorite settings. Some effect dialog boxes also provide graph controls for adjusting settings. As you adjust settings, you can use the Preview option to preview effects in real time.



For information on using specific effects, search for the effect name in Help or look in the index.

Using presets

Many of Adobe Audition's effects and other functions have presets that are available for easily storing and recalling your favorite settings. You can add and remove presets at any time.



Presets in the Amplify/Fade dialog box

To apply a preset:

Double-click the preset name. The settings defined by the preset are reflected in the dialog box.

To add a preset:

- 1 Adjust the effect settings as desired.
- 2 Click Add in the Presets area of the effect dialog box.
- 3 Enter a name for the preset, and click OK. Your new preset is added to the list of other presets, which is automatically sorted alphabetically.

To modify a preset:

- 1 Double-click the preset name, and adjust the settings as desired.
- 2 Click Add, enter the name of the current preset, and click OK.
- 3 Click OK when prompted to replace the preset.


To remove a preset

Select the preset, and click Del.

Using graph controls

Many of Adobe Audition's effects use graph controls for adjusting parameters. By adding and moving control points on the graph, you can tailor the effect to precisely meet your needs.

By default, graphs display straight lines between control points. However, some graphs provide a Splines or Spline Curves option for generating a curve between control points. Using spline curves lets you create smoother transitions between points.



Graph with straight lines between control points compared to graph with spline curves



When you use spline curves, the line won't travel directly through the control points. Instead, the points control the shape of the curve. To get the curve closer to a control point, click to create more control points near the point in question. The more control points there are clustered together, the closer the spline curve will be to those points.

To use graph controls:

Do any of the following:

- To add a control point to the graph, click in the grid at the location where you want to place the point.
- To enter the values for a control point numerically, right-click the point to bring up the edit box, or double-click the graph's curve.
- To move a point on the graph, drag it to a new location.
- To remove a point from the graph, drag it off the graph.

Note: When the pointer is located over a control point, you'll see it change from an arrow to a hand.

Previewing effects in Edit View

Many dialog boxes provide a Preview button for previewing effects in real time. This means that you can monitor the processed signal before applying the effect to the waveform. The preview feature updates in real time, meaning that changes you make to effect settings while in the dialog box for that effect become audible immediately, while the audio is playing.

Keep in mind that your system's performance affects the preview feature. On slower systems, some effects may tend to break up or skip during preview. In Multitrack View, the preview is not necessary, as effects are used nondestructively. Basically, every effect in the Multitrack View is in preview all the time. For more information on the differences between destructive and nondestructive editing, see [“About using Edit View and Multitrack View” on page 20](#).

In Edit View, you can add an optional preroll or postroll amount to the duration of the preview. This is especially useful when previewing effects for small ranges and marquee selections because it lets you hear how the in and out transitions are affected by the effects settings.

To preview effects in real time:

- 1 Click the Preview button to start playing the audio.
- 2 Adjust the effect settings as desired.
- 3 To compare the original audio to the processed audio, select and deselect the Bypass option. When the option is selected, you hear the original audio; when the option is deselected, you hear the processed audio.
- 4 When you're satisfied with the settings, click Stop.

To add a preroll and postroll duration to a preview:

- 1 In Edit View, right-click the Play button ▶ or the Play To End button ⏮ in the transport controls, and choose Preroll And Postroll Options.
- 2 In the Effects Preview section of the Preroll And Postroll Options dialog box, enter durations for the preroll and postroll, and click OK.

3 Do one of the following:

- Choose Effects > Enable Preroll And Postroll Preview.
- In an effects dialog box, select Enable Preroll And Postroll Preview. This option appears below the Presets. If a dialog box does not have Preset, the Enable Preroll And Postroll Preview option will not appear; however, you can still enable preroll and postroll preview by choosing Effects > Enable Preroll And Postroll Preview.

4 Preview an effect as described in the previous procedure.

Using plug-in effects

DirectX and VST plug-ins let you extend the already powerful effects at your disposal in Adobe Audition. Before you can start using plug-in effects, you must set them up in Adobe Audition. For DirectX effects, this process involves enabling the effects and then refreshing the effects list. For VST effects, you need to verify that Adobe Audition is scanning the directories where the effects are installed; then, you must refresh the effects list.

After that, using plug-in effects is as easy as using any other Adobe Audition effect. Just select an area to process, and choose the effect from the Effects > DirectX or Effects > VST menu (or from the Effects tab of the Organizer Window). Of course, you'll need to consult the documentation provided by the plug-in manufacturer for any help with its features.

Note: If Adobe Premiere Pro and Adobe Audition are installed on the same computer, Adobe Audition automatically displays the VST plug-ins that come with Adobe Premiere Pro.

To enable DirectX effects:

Do one of the following:

- In Edit View, choose Effects > Enable DirectX Effects.
- In Multitrack View, click the FX button in the track controls. In the Track Effects Rack dialog box, click Enable DirectX Effects, and then click OK.

This causes Adobe Audition to scan your system for DirectX plug-ins. After the plug-ins are activated, the Enable DirectX Effects option is removed from the menu and dialog box.

To set up directories for VST effects:

1 In Edit View, choose Effects > Add/Remove VST Directory.

The Add/Remove VST Directory lists the directories that Adobe Audition will scan for VST plug-ins when you choose Effects > Refresh Effects List.

2 Do either of the following:

- To add a new directory, click Add, locate or create the folder you want Adobe Audition to scan for VST plug-ins, and click OK.
- To remove a directory, select the directory and click Remove.

To refresh the effects list after installing new effects:

In Edit View, choose Effects > Refresh Effects List.

Setting up Adobe Audition

About setting up Adobe Audition

Setup tasks fall into several categories. Perhaps the most important is setting up the devices you want to use with Adobe Audition. If you have multiple sound cards, or a single card that has multiple inputs and outputs, you need to specify which devices you want to use for playback and recording. In addition, you can set up MIDI devices, external controllers, and ReWire connections for use with Adobe Audition. For more information on these tasks, see [“Setting up devices” on page 37](#).

Another category of setup tasks is customizing internal Adobe Audition preferences to best suit your needs. For example, you can change the appearance of the workspace, set buffer sizes to optimize performance, change the locations of temporary folders to better utilize disk space, and customize the wave and session displays. For more information on these tasks, see [“Setting Adobe Audition preferences” on page 43](#).

A final category of setup tasks is managing the size of temporary files. The size of temporary files is limited only by the amount of disk space that is available; however, when you're working with very large files (or when you have many files open at the same time), your disk space may run low. If this happens, you can delete temporary files you're not using, clear specific Undo items, and change the amount of reserve space. For more information on these tasks, see [“Managing temporary files” on page 53](#).

Setting up devices

You can use a wide range of devices with Adobe Audition. Sound card inputs let you bring audio signals into Adobe Audition through sources such as microphones, tape decks, and digital effects units. Sound card outputs let you monitor audio signals through sources such as speakers and headphones. MIDI ports let you connect Adobe Audition to MIDI keyboards and synthesizers. You can also synchronize Adobe Audition with ReWire applications and hardware or software components that support SMPTE/MTC timecode.

Designating which devices you want to use

The Device Order dialog box lets you designate which devices you want to use with Adobe Audition. When working in Edit View, you can designate one stereo output device to use for playback and one stereo input device to use for recording. When working in Multitrack View, you can assign different input and output devices to each audio track. However, before you can do this, you must specify which devices you plan to use and the order in which you want to view them.

If your audio system includes MIDI devices, you can also designate which MIDI input and output devices you want to use. For example, you can designate a MIDI keyboard to use for triggering commands and a MIDI synthesizer channel to use for playback. (See [“About using MIDI devices” on page 40](#).)



To designate the devices you want to use:

- 1 Choose Options > Device Order.
- 2 Click the tab for the type of device you want to designate: Playback, Recording, MIDI Output, or MIDI Input.
- 3 Move the devices you want to use into the Multitrack Device Preference Order list by selecting devices in the Unused list and clicking Use. Remove the devices you don't want to use by selecting devices in the Multitrack Device Preference Order list and clicking Remove.

Note: You can specify up to 16 stereo devices or 32 mono devices in the Multitrack Device Preference Order list.

- 4 Designate the device you want to use in Edit View by selecting the device and clicking Use in EV. [EV] appears after the device name.
- 5 Adjust the order of devices for use in Multitrack View by selecting a device and clicking Move Up or Move Down.

The first device in the list is the default device. This means that, by default, the first playback device is assigned as the output for all audio tracks in a session and the first recording device is assigned as the input for all audio tracks. Likewise, the first MIDI Out device is assigned as the output for all MIDI tracks. However, you can easily reassign the devices for a track. (See [“Using the Track Properties window” on page 197](#) and [“Importing and mapping MIDI files” on page 203](#).)

- 6 If desired, click a different tab to set up ordering for another type of device. When you are finished, click OK.



To quickly view or change the properties for a device, select the device and click Properties.

Setting properties for audio output devices

The Device Properties dialog box lets you specify Adobe Audition's parameters for playing back waveforms. If you have multiple sound cards, or a single card that has multiple audio outputs, you can customize the properties for each output.

To set properties for audio output devices:

- 1 Choose Options > Device Properties, and click the Wave Out tab.
- 2 Select a device from the list at the top of the dialog box.

The capabilities of the selected output device are shown in the Supported Formats table. A Yes or No indicates different combinations of sample rate and bit resolution. This table also shows what (if any) 32-bit formats the output device can handle, and whether it can accept the WDM driver extensible wave format.

- 3 Set any of the following properties. When you are finished, you can choose a different device to set up, or you can click OK to close the dialog box:

Order Displays the order of the device for use in Multitrack View. Click Change to open the Device Order dialog box and change the order of devices. (See [“Designating which devices you want to use” on page 37](#).)

Use This Device In Edit View Indicates that Adobe Audition will use the device to play waveforms in Edit View.

Limit Playback To Downsamples audio data for playback. Use this option to compensate for limitations imposed by your hardware. For example, if your sound card doesn't handle 32-bit audio correctly, you can have Adobe Audition limit the playback of 32-bit files to either 16-bit or 8-bit.

Send 32-bit Audio As Specifies how Adobe Audition sends 32-bit audio data to the output device. This option is not available if you select a Limit Playback To option. If the output device supports it, you can send 32-bit audio as 3-byte Packed PCM, 4-byte PCM, or 4-byte IEEE float.

Enable Dithering Activates dithering when playing back audio at a limited bit depth. If you deselect this option, Adobe Audition truncates the audio data instead. This means that bits that aren't used are simply chopped off and discarded. Enabling dithering is recommended when working with audio files that have a higher bit depth than your sound card supports. You can set the following options when dithering is enabled:

- **bits** specifies the number of bits to dither to. If you have a 20-bit sound card, for example, you will want to dither to 20 bits since any more bits will not be used by the card. Even for 16-bit-only sound cards, choosing to dither to 16-bit will improve the quality when playing back 32-bit audio.
- **p.d.f.** (probability distribution function) controls how the dithered noise is distributed away from the original audio sample value. Usually one of the Triangular p.d.f. functions is a wise choice, because it gives the best tradeoff between SNR, distortion, and noise modulation.
- **Shaping** specifies a noise shaping curve for moving noise to different frequencies. You can also specify that no noise shaping is used.

Setting properties for audio input devices

The Device Properties dialog box lets you specify Adobe Audition's parameters for recording waveforms. If you have multiple sound cards, or a single card that has multiple inputs, you can customize the properties for each audio input device.

To set properties for audio input devices:

- 1 Choose Options > Device Properties, and click the Wave In tab.
- 2 Select a device from the list at the top of the dialog box.

The capabilities of the selected recording device are shown in the Supported Formats table. A Yes or No indicates different combinations of sample rate and bit resolution.

- 3 Set any of the following properties. When you are finished, you can choose a different device to set up, or you can click OK to close the dialog box:

Order Displays the order of the device for use in Multitrack View. Click Change to open the Device Order dialog box and change the order of devices. (See [“Designating which devices you want to use” on page 37.](#))

Use This Device In Edit View Indicates that Adobe Audition will use the device to record waveforms in Edit View.

Get 32-bit Audio Specifies how the input device sends 32-bit audio data to Adobe Audition. If supported by the recording device, you can send 32-bit audio as 3-byte Packed PCM, 4-byte PCM, or 4-byte IEEE float.

Multitrack Latency Specifies the delay time (or *latency*) that the device introduces during recording. Many sound cards allow for monitoring input source signals with no latency. However, if you notice that tracks are out of sync, it is probably because one of the devices you used for recording introduced latency. Once you determine how out of sync a particular device gets, you can enter the number of milliseconds to delay a track's playback in relationship to all other tracks' playback to achieve synchronization.

Adjust To Zero-DC When Recording Removes any detected DC bias when recording.

About using MIDI devices

MIDI stands for Musical Instrument Digital Interface, and is a way of communicating performance information from one piece of software or hardware to another. This performance information can take the simple shape of a note instruction, as in E4, or it can transmit detailed information on things such as timing or sound patch data. Windows provides a way of transmitting MIDI information internally between programs, plus you can transmit MIDI information into and out of your computer to or from external devices (such as a MIDI Keyboard) through the MIDI port of a sound card, or other MIDI interface device.

You cannot record audio directly from a MIDI input device into Adobe Audition. In order to work with MIDI data in Adobe Audition, you must save the MIDI data to a file using a MIDI sequencing application, and then import the MIDI file into a session as a clip. Once you have MIDI clips in a session, you can map them to a specific MIDI output device and channel for playback. (See [“Working with MIDI tracks” on page 203.](#))

If you have a MIDI input device connected to your system's MIDI interface, you can use it to execute commands in Adobe Audition. For example, you can assign the Play command in Adobe Audition to the C4 note on your MIDI keyboard. This is called *MIDI triggering*. (See [“Using shortcuts” on page 21.](#))

You can also use your system's MIDI Out and In ports to send and receive SMPTE/MTC timecode. This process lets you synchronize Adobe Audition's Multitrack playback and recording with other hardware or software components that also support SMPTE/MTC. (See [“Setting up for SMPTE synchronization” on page 40](#) and [“Using sessions as SMPTE masters or slaves” on page 183.](#))

Setting up for SMPTE synchronization

You can use SMPTE (Society of Motion Picture and Television Engineers) timecode to synchronize Adobe Audition's transport controls with a MIDI sequencing application or an external hardware device, such as a videotape machine. (See [“Using sessions as SMPTE masters or slaves” on page 183.](#))

Adobe Audition sends and receives SMPTE timecode via MIDI timecode (MTC), which Windows transmits through your system's MIDI Out and MIDI In ports. MTC is a digital signal; to convert analog SMPTE timecode from a video or audio tape deck to digital MTC, you must use an appropriate MIDI interface.

To designate the devices with which you want to synchronize:

- 1 Choose Options > Device Properties.
- 2 Click the MIDI Out tab, and choose a device for SMPTE Output. This is the device to which Adobe Audition will send the MIDI timecode.
- 3 Click the MIDI In tab, and choose the device for SMPTE Slave Device. This is the device from which Adobe Audition will receive the MIDI timecode.
- 4 Click OK.

To set options for incoming SMPTE timecode:

- 1 If your MIDI interface supports sample-accurate synchronization, choose Options > Sample Accurate Sync.
- 2 Choose Options > Settings, and click the SMPTE tab.
- 3 Set the following options:

Lead Time Specifies the amount of time (in milliseconds) in which Adobe Audition establishes synchronization with incoming timecode. Lower settings (200 and lower) result in faster transport response but may prevent Adobe Audition from establishing synchronization. Settings of 500 to 1000 are sufficient on most systems.

Stopping Time Specifies the amount of time (in milliseconds) Adobe Audition will continue playing if it encounters a dropout in timecode.

Lag Time Specifies the number of samples between incoming timecode and outgoing audio data. This value accounts for discrepancies introduced by sound card buffers. The default value is 10 samples.

Slack Specifies the number of frames Adobe Audition can fall out of sync with timecode before either repositioning the current-time indicator to match the code or performing a full resynchronization. A setting of up to 2.5 frames is recommended, as incorrect timecode is usually corrected on the next frame sent. The default value is 1 frame.

Clock Drift Correction Time Specifies the number of samples to crossfade when making time corrections to chase audio to timecode. The default value is 200 samples.

Reposition Playback Cursor When Shuttling Readjusts the playback position if synchronization is off by the Slack value.

Full Re-Sync When Shuttling Performs a full re-synchronization if synchronization is off by the Slack value.

Setting up external controllers

You can use external controllers, such as the Mackie Control, when recording and mixing in Adobe Audition. These devices let you edit audio tracks using real knobs and automated faders, instead of your mouse and computer keyboard. The Device Properties dialog box lets you configure external controllers and specify a volume increment.

To set up external controllers:

- 1 Choose Options > Device Properties, and click the Ext. Controller tab.
- 2 Select the external controller you want to use, and specify a volume increment for the device.
- 3 Click Configure to set additional options for the device. These options are provided by the controller software. Refer to your controller documentation for more information.
- 4 Click OK.

Setting up ReWire connections

ReWire (a product of Propellerhead Software) is a technology for synchronizing audio applications. You can configure Adobe Audition to accept audio input from any ReWire-compatible application. When Adobe Audition is configured to accept ReWire input, it is referred to as a ReWire *host*. Applications that supply audio input are called ReWire *slaves* and the output channels they expose to the host are called devices.

To establish a ReWire connection, you first enable ReWire support in Adobe Audition and then activate a ReWire slave application and assign output from the slave to one or more Audition tracks. Adobe Audition serves as a ReWire host until you close the application. You can also manually disable ReWire support. For more information on using Adobe Audition as a ReWire host, see [“Working with ReWire tracks” on page 202](#).

Note: Before enabling ReWire in Adobe Audition you must close all other ReWire host and slave applications. After activating a slave application from within Adobe Audition, you will launch the application to establish the ReWire connection.

To establish a ReWire connection:

- 1 In Multitrack view, choose Options > Device Properties and select the ReWire tab.
- 2 Click Enable. The dialog box automatically populates with a list of installed ReWire slave applications.
- 3 Select the check box next to the application you want to activate as a slave.
- 4 Choose one of the following track assignment options:
 - Insert Summed Stereo Output Into First Available Track. All ReWire devices offer one summed stereo output. This option routes the summed stereo output into the first unoccupied track in the current session.
 - Insert All Outputs To Individual Tracks. ReWire devices may offer multiple channel outputs. This option routes each available ReWire output to its own track, starting with the first unoccupied track and following contiguously to additional unoccupied tracks.
 - Insert Outputs Manually Using Track Device Input Dialogs. Choose this option if you want to assign outputs manually by using the Input Device dialog box. (See [“Working with ReWire tracks” on page 202](#).)

5 Click Launch to launch the ReWire slave application and establish the ReWire connection. Adobe Audition assigns output from the ReWire slave to one or more tracks, as specified by the track assignment option you selected.

6 Open the session you want to work with in the ReWire slave application to make the audio available to Adobe Audition.

Note: Because only one ReWire host can be active at a time, you need to disable ReWire in Adobe Audition before enabling any other ReWire host application.

To disable ReWire support:

- 1 In Multitrack view, choose Options > Device Properties and select the ReWire Devices tab.
- 2 Click Disable, and then click OK.

Setting Adobe Audition preferences

The Settings dialog box lets you customize Adobe Audition's workspace, use of memory and hard disk space, spectral view, behavior when pasting, and other miscellaneous settings.

To use the Settings dialog box:

- 1 Choose Options > Settings.
- 2 Click a tab at the top of the dialog box to view the desired sets of options.
- 3 When you're finished setting options, click OK. To close the Settings dialog box without changing any options, click Cancel.

Once you click OK, most changes take effect immediately. If a change requires that you close and reopen Adobe Audition, you'll be prompted to do so. For example, you need to close and reopen Adobe Audition when you set up a different temporary folder.

General options

The General tab in the Settings dialog box provides options for adjusting mouse behavior in Adobe Audition, as well as parameters for auto-play, live update, auto-scroll, and more.

Force Spacebar To Always Trigger Play Forces the spacebar to always trigger playback regardless of which dockable window has focus.

Auto-Play On Command-Line Load Enables the ability to start Adobe Audition and play a file from the command line. For example, if you go to the Run command in the Windows Start menu and type "`c:\Program Files\Adobe\Audition 1.5\Audition.exe`" "`c:\Program Files\Adobe\Audition 1.5\Audition Theme\TalkBackVerb.cel`" at the command line, Adobe Audition will start and begin playing TalkBackVerb.cel.

Live Update During Recording Enables live waveform drawing while recording. On faster computers, you can have the waveform displayed in real time as audio is being recorded. However, if you find the recorded audio becoming choppy, disable this option.



In Edit View's Spectral View mode, and at lower spectral resolutions (around 256), you can perform a nice scrolling spectral plot while recording with this option on.

Auto-Scroll During Playback And Recording Enables scrolling of the waveform display in sync with playback. Auto-scrolling only takes affect when you are zoomed in on a portion of a waveform and play past the viewed portion.

Note: *The display refresh rate is directly related to the Total Buffer Size setting in the System tab of the Settings dialog box. A low buffer size (such as 1) results in a smooth scrolling display, where as a high buffer size (such as 8) results in a more choppy display. (See “System options” on page 45.)*

Upon A Manual Scroll/Zoom/Selection Change Determines auto-scrolling behavior when a manual scroll, a zoom, or a selection change occurs in Adobe Audition. You can abort auto-scrolling until the next time you play or record; resume auto-scrolling only when the play cursor enters the view; or resume auto-scrolling immediately. Choose the one that best suits your needs.

Custom Time Code Display Defines the number of frames per second (FPS) assigned to the Custom time format in the View > Display Time Format menu.

Restore Default Workspace Resets all window sizes and positions to Adobe Audition's default arrangement.

Ctrl Key Allows Dockable Windows To Dock Disables the Ctrl key from preventing a window to dock when moving the window around the work area.

Mouse Wheel Determines the amount to zoom in when rotating the mouse wheel found on Intellipoint-compatible pointing devices. Values from 10% to 80% work well. The higher the value, the further you'll zoom in when you roll the mouse wheel.

Time Selection Mouse Cursor Determines whether you want your mouse pointer to appear as an arrow or as an I-beam when it's over the waveform display.

Edit View Right-Clicks Determines the behavior for a right-click in the waveform display.

- **Popup Menu:** When right-clicking in the waveform display, a menu pops up if this option is selected. You can then Shift-click to extend a selection.
- **Extend Selection:** If you select this option, right-clicking in the Edit View's waveform display lets you extend the edge of a waveform selection instead of displaying the pop-up menu. To see the pop-up menu, hold down the Ctrl key as you right-click.

Default Selection Range Determines the amount of waveform data that automatically gets selected (if nothing is already highlighted) when you apply an effect.

- **View:** If this option is selected, the area that's automatically selected is limited to the area you can currently see on-screen.
- **Entire Wave:** When you choose this option, the entire waveform is automatically selected, even if you're only viewing a portion of it.



Double-clicking always selects the current view. Triple-clicking always selects the entire waveform.

Highlight After Paste Highlights the inserted selection when performing a Paste operation. Deselect this option to have the cursor placed at the end of the pasted selection instead.



Deselect this option for easier multiple pastes, one after the other.

System options

The System tab in the Settings dialog box provides options for configuring how Adobe Audition interacts with your system.

Edit View Play/Record Buffer Determines the buffer size (in seconds) to be used when sending data to and from your sound card when playing back or recording in Edit View. Different sound card devices may require different memory buffer settings. The default settings should work fine for most sound cards, but if you hear choppiness (skips or dropouts) in recording or playback, you may need to adjust the buffer size or number of buffers used. For example, if you experience breakups in your audio, or you can't stop a recording in progress, increase the buffer size.

Use the two fields in the Edit View Play/Record Buffer area to reserve more memory for recording and playback by entering a higher buffer size, both in seconds and a number of buffers.



Keep in mind that while a greater buffer size will allow for increased multitasking when audio is being played, it does so at the expense of taking more of your computer's memory.

Wave Cache Determines the amount of memory that Adobe Audition reserves for processing data. Recommended cache sizes are from 8192 to 32768 KB (8192 KB is the default).

Select Use System's Cache to let Windows handle all disk caching. Keep in mind that Adobe Audition usually handles caching better than Windows can. However, this option reserves the least amount of memory, so it may be desired for systems with low amounts of RAM.

EV Preview Buffer Determines the minimum buffer size used when sending data to your sound card for the real-time Preview feature found in many effect dialog boxes. The default value is 250 milliseconds.

Different sound card ports may require different memory buffer settings. If you hear choppiness (skips or dropouts) when you use the Preview feature, try adjusting the buffer size used. (Choppiness can be caused by insufficient processing power as well.) Keep in mind that a larger Minimum Preview Buffer Size requires more computer memory.

Use Sound Card Positioning Info Allows Adobe Audition to query the sound card for the actual location and sync up the cursor with audio. This option is useful if a sound card doesn't play or record at 44,100 Hz (some sound cards, for example, work at 44,050 Hz or 44,130 Hz). Leave this option unselected unless the cursor is out-of-sync with the audio.

CD Device Options Specifies the SCSI interface used by your CD device: ASPI (Advanced SCSI Programming Interface) or SPTI (SCSI Pass Through Interface).

Temporary Folders Specifies the folders in which you want Adobe Audition to store temporary files. Adobe Audition creates temporary files for use when performing edits on your audio. All temporary files begin with CEP and have the .tmp extension. The rest of the filename is chosen at random when the file is created. If there are no copies of Adobe Audition running, none of these files should be present, since Adobe Audition normally deletes temporary files when it exits. However, these files can be left behind in extreme circumstances if Adobe Audition crashes, or if Windows unexpectedly quits while Adobe Audition is active. As long as Adobe Audition isn't running, you can safely delete these files. You can also use the Manage Temporary Folder Reserve Space to delete temporary files you aren't using while Adobe Audition is running. (See [“Managing temporary files” on page 53.](#))

Important: *You need to have enough space available in these folders to accommodate the total size of all the audio files you wish to edit simultaneously.*

Use the reserve free fields to specify an amount to leave available for headroom purposes for both the primary and secondary temporary folders.

- **Temp Folder:** Specifies Adobe Audition's main temporary folder. Ideally it should be on your fastest hard drive.
- **Secondary Temp:** Specifies Adobe Audition's secondary temporary folder. For best results, this should be on a different physical hard drive than the primary temp folder. This is especially true when recording more than one track at a time in Multitrack View, because odd track recordings go to one temp folder while the even tracks are recorded to the other temp folder, dividing the workload.

Note: *Providing you have enough free space on the drive that holds the primary temporary folder, Adobe Audition will work just fine if no Secondary Temp folder is specified.*

Undo Specifies options for Adobe Audition's Undo feature, which lets you revert back to your last edit with a keyboard shortcut (Ctrl+Z), menu command, or toolbar button.

- **Enable Undo:** Activates the Undo function. Because Undo requires extra disk space for its temporary files and time to save them before processing, you may sometimes want to turn this feature off.
- **Levels (minimum):** Specifies the fewest number of Undo levels.
- **Purge Undo:** Deletes all of Adobe Audition's Undo files. This frees up disk space, but ends your ability to revert to previous edits.

Delete Clipboard Files On Exit Deletes Adobe Audition clipboard files when you exit. In general, leave this option enabled: Usually, after you finish with an Adobe Audition session, these clipboard files are no longer needed and just take up valuable hard disk space.

Deselect this option to retain Adobe Audition's clipboard files on your hard drive after you exit the program.

Force Complete Flush Before Saving Disables the quick save feature, in which Adobe Audition quickly saves files that contain only minor modifications. If you enable this option and force a flush before saving, Adobe Audition saves all files by making a backup copy of the file internally and then writing the entire file back.

This option is disabled by default. When enabled, it considerably increases the save time for large files. It is intended for use only if you have trouble saving back to the same filename or you have a problem with Adobe Audition's quick save feature.

Colors options

The Colors tab in the Settings dialog box provides options for changing Adobe Audition's color scheme.

Color Presets Lists color scheme presets that come with the program as well as those you've created yourself. To choose one, select it from the list. The currently selected color scheme is displayed in the Example window.

Save As Saves the currently selected color scheme as a preset.

Delete Deletes the currently highlighted color scheme preset.

Waveform Tab Lists all of Adobe Audition's waveform elements to which you can assign custom colors. Choose an item from the list and click the Change Color button to change the color.

To adjust the appearance of the selected (highlighted) portions of waveforms and blocks, select a Selection option:

- **Transparency:** Drag the slider or enter a value to adjust the transparent value (in percentage) of a selection; 0 is no transparency and 100 is maximum transparency.
- **Invert:** Select to set the selection colors to the inverse of the nonselected colors.

Spectral Tab Lists the display elements for Adobe Audition's spectral display. Select an item from the list and click the Change Color button to adjust the element's color.

For Spectrum, choose one of these options:

- **Reverse Direction:** Inverts the normal colors of the spectrum display, similar to an Invert or Negative command in a photo editor.
- **Gamma:** Adjusts the overall brightness of the Spectral View. Positive numbers make the display brighter, while negative numbers darken the display. This setting works just like the Gamma function in many image editors.

To adjust the appearance of the selected (highlighted) portions of waveforms in spectral display, choose a Selection option.

- **Transparency:** Drag the slider or enter a value to adjust the transparent value (in percentage) of a selection; 0 is no transparency, and 100 is maximum transparency.
- **Invert:** Select to set the selection colors to the inverse of the nonselected colors.

Controls Tab Lists the Adobe Audition control elements for which you can change colors. Select an item from the list and click the Change Color button to adjust the element's color.

Select Segmented Progress Bar to make the progress bar segmented instead of solid. The progress bar appears when you apply an effect, or open or save large waveforms.

Select White Progress Background to make the background of the progress bar white.

For Dockable Windows, select one of the following options:

- **Use System 3D Color:** Select to make dockable windows use your system's 3D color. This is the color Windows uses to render most windows on your system.
- **Use Darkened System 3D Color:** Select to make dockable windows use the darkened version of your system's 3D color.
- **Use Specified 3D Color:** Select to make Adobe Audition's dockable windows use the 3D color you specify.

To change the 3D color, select Dockable Windows 3D Color in the controls list, and then click the Change Color button to select a new color.

Display options

The Display tab in the Settings dialog box provides options for adjusting Adobe Audition's Spectral View and Waveform View modes.

Windowing Function Determines the method Adobe Audition uses to segment the spectral data before it displays it. The segments (windows) are listed in order from the narrowest frequency band/most noise to the widest frequency band/least noise. Blackmann or Blackmann-Harris are good choices.

Resolution Specifies the number of vertical bands used in drawing frequencies. Keep in mind that the larger this number, the longer it will take for Adobe Audition to render the spectral display. Performance will vary based on the speed of your computer.

Window Width Specifies the width of the window (or frame size) used in plotting the spectral data, where 100% is a frame size of the FFT (Fast Fourier Transform) size. Window Width basically lets you increase time resolution at the expense of some frequency resolution. So the display will become more accurate along the timeline (left and right) and less accurate along the frequency scale (up and down) as the window width decreases. The default setting is 75%, but you should lower the value (50 to 75% works best) if you want to increase the resolution horizontally—for example, to find out exactly where a certain frequency starts.

Plot Style Specifies a style for plotting frequencies:

- **Logarithmic Energy Plot:** In this mode, colors change with the decibel value of the energy at any particular time and frequency. In this mode, you can see more details in the very quiet ranges, especially if the Range is quite high (above 150 dB). Use the Range value to adjust the sensitivity in plotting frequencies.
- **Linear Energy Plot:** When selected, colors are chosen based on percentage of maximum amplitude instead of decibel amplitude. Linear Energy Plot can be useful for viewing the general overview of a signal without getting bogged down by detail at much quieter levels. You can adjust the Scaling factor to highlight audio of different intensities.

Show Cue And Range Lines Displays cue marker and range lines in the waveform display. Cue marker and range entries in the Cue List appear with vertical dotted lines overlaying the audio, connecting the arrows from the top to the bottom of the display.

Show Grid Lines Displays grid lines in the waveform display. The grid lines mark off time on the horizontal x-axis and amplitude on the vertical y-axis.

Show Center Lines Displays center lines in the waveform display. The center lines represent zero amplitude of the waveform's right and left channels.

Show Boundary Lines Displays boundary lines in the waveform display. Boundary lines are the horizontal lines that visually indicate where the waveform's amplitude approaches or exceeds the clipping level. The value in the Display Boundary Lines At option specifies the amplitude at which the boundary lines appear.

Peak Files Specifies options for peak (.pk) files, which Adobe Audition uses to store information about how to display WAV files. Peak files make file opening almost instantaneous by greatly reducing the time it takes to draw the waveform (especially with larger files).

- **Peaks Cache:** Determines the number of samples per block to be used when storing peak files. Larger values reduce the RAM requirement for large files at the expense of slightly slower drawing at some zoom levels. If RAM is an issue on your system, and you're working with very large files (several hundred megabytes or more in size), consider increasing the Peaks Cache to 1024 or even 1536 or 2048.
- **Save Peak Cache Files:** Specifies that peak files are saved with all .wav files (in the same folder) with the extension .pk following the original audio filename.
- **Rebuild Wave Display Now:** Click to rescan the current file for sample amplitudes and redraw the waveform.

Data options

The Data tab in the Settings dialog box provides options for controlling how Adobe Audition handles audio data.

Embed Project Link Data For Edit Original Functionality Links session files with exported mixdown files. Once these files are linked, you can select a mixdown file in Adobe Premiere Pro or After Effects, and then open and remix the related session in Adobe Audition's Multitrack View.

Auto-Convert All Data To 32-Bit Upon Opening Converts all 8-bit and 16-bit data to 32-bit when a file is opened, and all subsequent operations will keep the data in the 32-bit realm.

Interpret 32-Bit PCM .wav Files As 16.8 Float Causes this version of Adobe Audition to be compatible with previous versions when it comes to handling 32-bit PCM .wav files.

Dither Transform Results (increases dynamic range) Enables dithering when processing effects such as FFT Filter or Amplify. Most processing done by Adobe Audition uses arithmetic greater than 16-bit, with the results converted back to 16-bit when complete. During this conversion, dithering provides a higher dynamic range and cleaner results, with less distortions and negative artifacts. With dithering, you get almost 24-bit sample performance in only 16 bits, as the dynamic range is increased by another 10 dB or so, allowing signals as quiet as -105 dB.

If this option is disabled, the results are truncated to 16 bits when converting back, thus losing the more subtle information.

When enabled, the addition of dither retains this subtle information. The drawback is that with each operation a small amount of white noise is added at the quietest volume level. However, the trade-off between using dither (thus adding noise) and truncating the data (thus creating artifacts and correlated quantization noise) generally favor using dither, so it's best to leave this option enabled.

Use Symmetric Dithering Enables symmetric dithering. In most cases, it's best to leave this option selected. If unselected, a DC offset of one-half sample is added each time data is dithered. Symmetric dithering has just as many samples added above zero as below zero. By contrast, nonsymmetric dithering just toggles between 0 and 1. Sometimes in a final dither, this may be desired to reduce the bit range of the dither. However, both methods produce identical audible results in every respect.

Smooth Delete And Cut Boundaries Smooths Cut and Delete operations at the splicing point, preventing audible clicks at these locations.

Smooth All Edit Boundaries By Crossfading Automatically applies a crossfade to the starting and ending boundaries of the selection. This option smooths any abrupt transitions at these endpoints, thus preventing audible clicks when filtering small portions of audio. You can enter a value (in milliseconds) in the crossfade time box to specify the crossfade duration to be applied.

Auto-Convert Settings For Paste When pasting different sample formats, Adobe Audition uses these settings when auto-converting the clipboard to the current sample format. Valid settings range from 30 to 1000.

- **Downsampling Quality Level:** Enter a value (30 to 1000) for downsampling quality. Higher values retain more high frequencies while still preventing the aliasing of higher frequencies to lower ones. A lower quality setting requires less processing time, but results in certain high frequencies being rolled off, leading to muffled-sounding audio. Because the filter's cutoff slope is much steeper at higher quality settings, the chance of ringing at high frequencies is greater. Usually values between 80 and 400 do a great job for most conversion needs. The default value is 80.
- **Pre-Filter:** To prevent any chance of aliasing, the pre-filter on downsampling removes all frequencies above the Nyquist limit, thus keeping them from generating false frequencies at the low end of the spectrum. In general, select this option for best results.
- **Upsampling Quality Level:** Enter a value (30 to 1000) for upsampling quality. Higher values retain more high frequencies while still preventing the aliasing of higher frequencies to lower ones. A lower quality setting requires less processing time but results in certain high frequencies being rolled off, leading to muffled-sounding audio. Because the filter's cutoff slope is much steeper at higher quality settings, the chance of ringing at high frequencies is greater. Usually values between 100 and 400 do a great job for most conversion needs. The default value is 120.



You should use a higher value whenever you downsample from a high sample rate to a low rate. For upsampling, a lower value produces quality almost identical to a higher value. The difference lies in the larger phase shift that exists at higher frequencies, but since the phase shift is completely linear, it's very difficult to notice. Downsampling, at even the lowest values, generally doesn't introduce any undesired noisy artifacts. Instead, the sound might be slightly muffled because of the increased high-end filtering.

- **Post-Filter:** To prevent any chance of aliasing, the post-filter on upsampling removes all frequencies above the Nyquist limit, thus keeping them from generating false frequencies at the low end of the spectrum. In general, select this option for best results.

Dither Amount For Saving 32-Bit Data To 16-Bit Files Enables dithering when pasting 32-bit audio to 16-bit. The default value of 1 (bit) enables dithering, while a value of 0 disables dithering. For semi-dithering, choose a value of 0.5.

With dithering, you get almost 24-bit sample performance in only 16 bits, as the dynamic range is increased by another 10 dB or so. This allows signals as quiet as -105 dB.

Allow For Partially Processed Data After Canceling Effect Determines what happens after you click the Cancel button while in the middle of applying an effect to a waveform. When selected, Adobe Audition leaves the effect applied to all data processed up until the point you clicked Cancel. When deselected, Adobe Audition automatically removes the effect on already processed data when you click Cancel.

Multitrack options

The Multitrack tab in the Settings dialog box provides options that let you optimize performance during recording, playback, and mixdown.

Playback Buffer Size Determines the buffer size (in seconds) used when sending data to your sound card when playing back a multitrack session. Different sound card drivers may require different memory buffer size settings. Adobe Audition's default settings should work fine for most sound cards. If you hear choppiness (skips or dropouts) in multitrack playback, adjust the buffer size. (Choppiness in multitrack playback can also be attributed to the background mixing process not being far enough ahead). A larger buffer size requires more computer memory. The default setting is 1.

Playback Buffers Specifies the number of buffers Adobe Audition uses for playback in the multitrack environment. If you experience break-up in your audio, try reducing the number of buffers. Increasing this number might also be helpful for some configurations. The default setting is 10.

Recording Buffer Size Reserves memory for recording in a multitrack session by entering a buffer size (in seconds). Different sound card drivers may require different memory buffer size settings. Adobe Audition's default settings should work fine for most sound cards. If you experience dropouts while recording in multitrack (especially when playback seems fine), try increasing this setting. (First, be sure the background mixing process is sufficiently complete when you go to record as this may cause the same symptom.) A larger buffer size requires more computer memory. The default setting is 2 seconds.

Recording Buffers Specifies the number of buffers used for recording in the multitrack environment. If you experience break-ups in your audio, try reducing the number of buffers. Increasing this number may also help for some configurations. The default setting is 10.

Background Mixing Priority Specifies the priority level of the background mixing process in a multitrack session. Lower values indicate a higher level of priority above other system events. You can use fractional numbers (such as 0.8). The default setting is 2.

Open Order Determines the order in which Adobe Audition opens a sound card's playback (in) and record (out) ports for use in the multitrack environment. This order is relevant only for older sound cards that don't support full-duplex capability.

Start Order Determines the order in which Adobe Audition starts a sound card's playback (in) and record (out) ports for use in the multitrack environment. This order is relevant only for older sound cards that don't support full-duplex capability.

Correct For Drift In Recordings Synchronizes the master audio playback device (generally, the first Out device listed in the session—the one on Track 1) and the record device of the waveform being recorded. If the true sample rates on the cards differ enough that the recording would have drifted out of sync with the original if both were played back at exactly the same sample rate, then the recording is corrected by resampling to make it the proper length. This option only works with new record tracks, not with recording on top of existing waveforms, or punch-ins.

Note: *On sound cards that support sample accurate devices (that is, synchronized device starting, and all devices keyed off of the same clock) you don't need to select this option. This option allows for some measure of near sample-accurate synchronization across different sound cards, or when using with a single sound card that doesn't use the same clock for playback and recording (which is common in consumer and other low-end sound cards).*

Correct For Start Sync In Recordings Compares the exact true time that the record device started with the time the master playback device started. If different, the recorded block's position is adjusted so the recording starts in perfect sync with the playback. This option only works with new record tracks, not with recording on top of existing waveforms, or punch-ins.

If this option is enabled, and you do a loopback test (by connecting the audio Out to the audio In and recording some ticks) and each recording is still a fixed amount out of sync, then you can adjust for this by entering this amount (in milliseconds) in the Latency field of Options > Device Properties for the recording device being used. To compute milliseconds, look at the difference in samples, multiply by 1000, and then divide by the sample rate. For instance, if the recording consistently appears 27 samples ahead of the playback, the latency would be $27 \times 1000 / 44,100$, or about 0.61 milliseconds. (The reason for the milliseconds format and not samples is because at various sample rates this latency will be different in terms of samples, but will be the same in terms of milliseconds.)

Note: *On sound cards that support sample accurate devices (that is, synchronized device starting, and all devices keyed off of the same clock) you don't need to select this option. This option allows for some measure of near sample-accurate synchronization across different sound cards or a situation where a single sound card uses different clocks for playback and recording. (This situation is common for consumer and other low-end sound cards.)*

Delete Old Takes After Merging Automatically deletes any unused takes created during a punch-in when you select a take. If you don't select this option, unused takes remain available to the Session (in the Insert menu) and occupy hard drive space.

Crossfade Time Determines the amount of time (in milliseconds) over which crossfading occurs when a take created using punch-in is merged back into the surrounding waveform.

Mixdowns Determines the bit-resolution that is used when performing a mixdown. Regardless of the session format (16-bit or 32-bit), you can generate mixdowns at either 16-bit or 32-bit quality with this option. The default is 16-bit. Click Dithering Option to specify how to dither the 16-bit mixdown.

Track Record Specifies how waveforms are created when recording directly into the Multitrack View: as mono or stereo, and as 16-bit or 32-bit.

Pre-Mixing Determines the bit size used for the background mixing process. Best quality is achieved by leaving this at the default 32-bit setting. However, if you're using multiple sound cards, it may be advantageous and faster to choose 16-bit for pre-mixing, as less data will be transferred across the hard drives. For single output device situations, or faster hard drives, 32-bit is better as it provides optimization at mixdown.

Panning Mode Specifies the method used for panning waveforms in a multitrack session.

- **L/R Cut (log):** Pans left by reducing the volume of the right channel, and pans right by reducing the left channel volume. The channel being panned to doesn't increase in volume as panning gets closer to 100%.
- **Equal-power Sine:** Pans left and right channels with equal power, so a hard pan left will contain the same loudness as both channels together. This results in an increase of 3 dB RMS on the channel being panned to when at 100%.

Note: Because panning can actually make one channel louder than the original waveform, audible clipping can occur in 16-bit sessions. To avoid this, work in the 32-bit realm if you're using the Equal-power Sine panning method.

Auto Zero-Cross Edits Automatically adjusts the beginning and end points of all Cut, Copy, and Paste-type edits to the nearest place where the waveform crosses the center line (zero amplitude point).

If the amplitudes aren't lined up on both sides of the selection, the endpoints are at different amplitudes. This often results in an audible pop or click at that point.

Smooth Auto-Scrolling During Playback Enables smooth scrolling when playing back audio in Multitrack View. By default Adobe Audition uses a paging method of scrolling in Multitrack View instead of the smooth scrolling technique used in Edit View. This saves on system resources.

Save Locked Track Files After Closing Sessions Saves the temporary files associated with locked tracks. When you reopen the session, Adobe Audition uses the temporary file instead of mixing down the locked tracks.

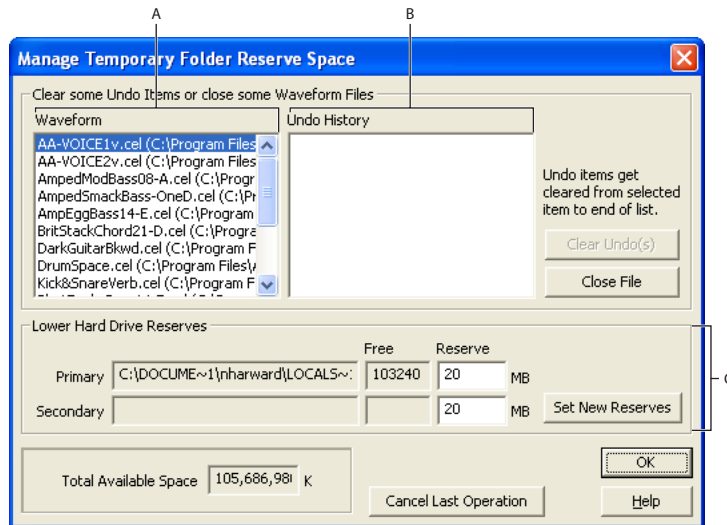
SMPTE options

The SMPTE tab in the Settings dialog box provides options for adjusting the settings for incoming SMPTE timecode. For more information, see ["Setting up for SMPTE synchronization" on page 40](#).

Managing temporary files

When you edit a file, Adobe Audition converts the audio data into an internal, temporary waveform. This process allows for quicker editing, better handling of large files, and the ability to undo changes. You can specify the folders where you want Adobe Audition to save temporary files and customize Undo options in the System tab of the Settings dialog box. (See ["Setting Adobe Audition preferences" on page 43](#).)

One advantage to using temporary files is virtually unlimited waveform sizes, since the maximum waveform size depends only on the size of your hard drive. The drawback, of course, is that the temporary file can get extremely large, potentially preventing you from being able to save a masterpiece on the same drive. If you notice long delays between edits or stuttering sounds on playback, you may be running out of free disk space in which to save the temporary file. In this case, you can use the Manage Temporary Folder Reserve Space dialog box to delete temporary files you're not using, clear specific Undo items, and change the amount of reserve space. This dialog box automatically appears when available hard drive space nears zero kilobytes.



Manage Temporary Folder Reserve Space dialog box

- A.** Open waveforms **B.** Undo items for the selected waveform
C. Location of primary and secondary temporary folders



Use the Status Bar to monitor the amount of free disk space. (See [“Using the status bar” on page 27](#)).

Adobe Audition doesn't create a temporary file for a waveform until you edit the waveform. However, you can force Adobe Audition to create a temporary file by using the Flush Virtual File command. This is useful when you need to use a waveform simultaneously in Adobe Audition and another application.

To manage temporary folder reserve space:

- 1 Choose File > Manage Temporary Folder Reserve Space.
- 2 Do any of the following:
 - To close a temporary file you're no longer using, select the file in the Waveform list, and click Close File. (The currently active waveform can't be closed this way, however.)
 - To clear Undo items for a file, select the file in the Waveform list. The Undo History list displays the actions that are currently being retained on your system and the amount of hard drive space each instance consumes. Select an item and click Clear Undo(s). All items at the selected level and below are removed.
 - To change the amount of space you want to keep free on the drives where the temporary files reside, enter a value in the Reserve text box, and click Set New Reserves.
 - To stop any action in progress, such as the application of an effect or any other edit, click Cancel Last Operation. This option is useful only if the dialog box automatically appeared because you ran out of storage space.



If Adobe Audition crashes, there may be a temp file (CEPx*.tmp) in your temporary folder that you should manually delete.

To force Adobe Audition to create a temporary file for the current waveform:

In Edit View, choose File > Flush Virtual File.

Importing, Recording, and Playing Audio

Opening audio files and multitrack sessions



Both Edit View and Multitrack View provide a variety of methods for opening files. In Edit View, you can open audio files; in Multitrack View, you can open session files.

Opening audio files in Edit View

In Edit View, you can open audio from a variety of audio file formats, including MP3, WAV, and AIFF. For more information on supported file formats, see [“Choosing an audio file format” on page 240](#).

If desired, you can change the sample type of the audio when you import it or append the audio to the end of the current waveform. Whichever method you choose for opening files, Adobe Audition provides options that let you preview the contents of files before you open them.

To open an audio file:

- 1 In Edit View, choose File > Open. Alternatively, click the Open button  in the toolbar or the Import button  in the Files tab of the Organizer window.
- 2 Locate and select the file you want to open. To select multiple, adjacent files, click the first file and Shift-click the last. To select multiple, nonadjacent files, Ctrl-click them.

Note: If you don't see the name of the file you want, choose All Supported Media from the Files Of Type menu. If you still don't see the file, it might be stored in a format that Adobe Audition can't read.

- 3 Click Open.

To preview the contents of a selected file:

Do any of the following:

- Click Play to listen to the file once.
- Select Loop to repeat the file until you click Stop.
- Click Auto Play to play files automatically when you select them.



To append an audio file to the current waveform:

- 1 In Edit View, choose File > Open Append.

If the new audio has a different sample rate, resolution, or channel type than the current waveform, Adobe Audition converts it to match the current waveform. For the best results, append files that have the same sample rate as the waveform.

- 2 Locate and select the file you want to open. To select multiple, adjacent files, click the first file and then Shift-click the last. To select multiple, nonadjacent files, Ctrl-click them.
- 3 Click Open.

To convert audio to a different sample rate, resolution, or channel type during import:

- 1 In Edit View, choose File > Open As.
- 2 Locate and select the file you want to open, and click Open.
- 3 Set the desired options in the Open File(s) As dialog box, and click OK:

Sample Rate Determines how many frequencies can be encoded in the audio signal. (Higher sampling rates mean wider bandwidth.) For more information, see [“About sample rates” on page 93](#).

Channels Determines if the waveform is mono or stereo. Select Mono to create a waveform with just one channel of audio information. This option works well for a voice-only recording. Select Stereo to create a two-channel waveform with separate right and left channels. This option is usually best for a music recording. Because they contain twice as much data, stereo waveforms consume twice the storage space of mono waveforms.

Resolution Determines the number of unique amplitude levels Adobe Audition can use to represent a sound. The 32-bit level is best while you work in Adobe Audition, and you convert down for output if necessary.

Note: Older sound cards might not be able to play 32-bit files properly. To check the capabilities of your sound card, choose Options > Device Properties. If your sound card doesn't support 32-bit files, you can convert the files to a lower bit rate (such as 16-bit) for playback.


Opening session files in Multitrack View

Session files contain no audio data themselves. Instead, they are small files that point to other audio files on the hard drive. A session file keeps track of what files are a part of the session, where they go in the multitrack, what envelopes and effects are applied to the tracks, and so on. For more information on creating session files, see [“Creating new sessions” on page 180](#).

In Multitrack View, you can open individual session files and you can append one session to another to quickly build elaborate compositions with shared themes. When you append sessions, appended tracks appear below current tracks. For example, if the current session uses tracks 1-4, Adobe Audition appends tracks 5 and greater and places them at the beginning of the timeline. If desired, you can then move clips in appended tracks to a new position. (See [“Selecting and moving clips” on page 188](#).)

Note: You can append a session only if it uses the same sample rate and bit depth as the current session. The sample rate and bit depth of the current session is displayed in the status bar.

To open a session file:

- 1 In Multitrack View, choose File > Open Session. Alternatively, click the Open button  in the toolbar.
- 2 Locate and select the file you want to open, and click Open.


To append a session file to the end of the current session:

- 1 In Multitrack View, choose File > Append To Session.
- 2 Locate and select the file you want to open, and click Open.

Inserting audio files into multitrack sessions

When you insert an audio file in Multitrack View, the file becomes an audio clip on the selected track. For more information about audio clips, see [“Working with clips” on page 187](#).

To insert an audio file into a multitrack session:

- 1 In Multitrack View, position the current-time indicator at the desired insertion point.
- 2 Select the desired track.
- 3 Do one of the following:
 - Choose Insert > Audio, select the audio file, and click Open. To preview the contents of a selected file, click Play to listen to the file once, or click Auto Play to play the file automatically when you select it. Select Loop to repeat the file until you click Stop.
 - Choose Insert, and select the name of a recently opened waveform from the submenu.
 - Choose Insert > File/Cue List. A window appears that lists all of the files that are currently open in Edit View. If a file has cues in it, a plus sign (+) appears next to its name to let you expand that file and see all the cue ranges in it. Click the file or cue you want to insert. Alternatively, drag the file or cue into the track display.
 - Select one or more files in the Files tab of the Organizer window, and click the Insert Into Multitrack button . If you select multiple files, each is inserted into a separate track. This method lets you insert a file into a session without leaving Edit View. (See [“Organizing files” on page 30](#).)

Note: If the audio file is longer than the space available on the selected track, Adobe Audition inserts the new clip on the nearest empty track.

Importing audio from CD

If you want to import audio into Adobe Audition from a CD, you can digitally extract it or record it internally. Digital extraction is the recommended method because it produces higher-quality audio than internal recording. Only use internal recording if your CD-ROM drive doesn't support digital extraction.

Extracting tracks from CDs

If your computer's CD-ROM drive supports audio digital extraction (also known as *ripping*), you can extract tracks from audio CDs. Once the audio is in Adobe Audition, you can edit it like any other waveform. Of course, if the CD is a typical read-only compact disc, you won't be able to save those changes back to CD. Instead, save modified CD tracks to a hard disk or burn them onto a new CD.

Adobe Audition provides two methods for ripping tracks from CDs: using the Open command and using the Extract Audio From CD command. Using the Open command is the quickest method and is preferred for ripping entire tracks. Using the Extract Audio From CD command gives you more control, such as the abilities to rip partial tracks and specify the ripping process used.

To extract tracks from a CD by using the Open command:

- 1 Place an audio CD in the computer's CD-ROM drive.
- 2 In Edit View, choose File > Open.
- 3 Choose CD Digital Audio (*.cda) as the file type, and navigate to the computer's CD-ROM drive.
- 4 Select the tracks you want to rip, and click Open.

To extract tracks from a CD by using the Extract Audio From CD command:

- 1 Place an audio CD in the computer's CD-ROM drive.
- 2 In Edit View or CD Project View, choose File > Extract Audio From CD.
- 3 For Device, choose the drive that contains the audio CD.
- 4 For Source Selection, do one of the following:
 - Select Track to extract one or more complete CD tracks. A list of all tracks on the CD appears, along with their lengths stated in Min:Sec:Frame format. (Each second of CD audio has 75 frames.)
 - Select Time to extract part of a track or a segment of audio that spans multiple tracks. Enter the beginning frame in the Start box, and the total number of frames you wish to extract in the Length box. (Each second of CD audio has 75 frames.) The actual start and length times appear in Min:Sec:Frame format above their respective boxes. The Range bar provides a graphical representation of how much audio will be extracted and where the audio appears within the CD. However, if you select only a short bit of audio to extract, you might not see any change in the Range bar.



The Time option is great for pulling hidden tracks from CDs, as well as for joining tracks that have been broken up by track indexes (such as performance track CDs and live albums).

- 5 For Interface Option, choose Generic Win32 or ASPI/SPTI. In most cases, ASPI/SPTI is the best choice. Select Generic Win32 only if the ASPI/SPTI option doesn't produce satisfactory results. The Generic Win32 option causes the Extract Audio From CD feature to use Input/Output control codes instead of SCSI commands.

For more information on ASPI/SPTI options, see [“Extract Audio From CD options” on page 60](#).

6 For Error Correction, CDDA Accurate is automatically selected if the CD-ROM drive has built-in ripping error correction. For these types of drives, no error correction is needed, so you won't be able to select any options from this part of the Extract Audio From CD dialog box.

However, if your drive isn't CDDA Accurate, you have access to No Correction and Jitter Correction options. No Correction, as you'd expect, means that no error correction will be performed. Jitter Correction compensates for data reading problems that older drives might have.

7 To listen to the selected tracks before extracting them, click Preview.

8 To save the settings for future use, save a preset. (See [“Using presets” on page 33.](#))

9 After you finish setting options, click OK.

Extract Audio From CD options

If you select ASPI/SPTI in the Extract Audio from CD dialog box, set the following options as desired:

Read Method Lets you choose the way Adobe Audition reads CD audio. Several methods are provided, many of them developed before the SCSI 3 specifications were published. (The SCSI 2 specs don't accommodate CD ripping.)

- MMC – Read CD is a SCSI 3-specific setting, and it works with most all recent drives. If you have a newer CD-ROM drive, try this setting first.
- SBC – Read10 is a standard SCSI read setting that uses a 10-byte SRB (SCSI Request Block). All SCSI devices are required to support this setting.
- SBC – Read6 is a standard SCSI read setting that uses a 6-byte SRB (SCSI Request Block). Many SCSI devices support this setting, but because it's optional, not all do.
- Plextor (D8) sends the D8 SCSI Op Code to the CD-ROM drive. Use this setting with older Plextor CD-ROM drives.
- D5 sends the D5 SCSI Op Code to the CD-ROM drive.
- NEC works with older NEC CD-ROM drives.

CD Speed Lists all extraction speeds that the selected CD-ROM drive supports and lets you specify the speed you want to use. The Max (Maximum) Speed option usually produces satisfactory results, but if it produces errors, specify a slower speed.

Buffer Size Specifies how much data Adobe Audition calls into the CD Extraction module to fetch, therefore determining how much data is pulled from the CD in each call to the read command. The default is 16 KB, but you can experiment with other sizes (which range all the way to the highest buffer size the CD-ROM drive supports). Although higher sizes mean faster ripping, they *could* introduce errors into the ripped file.

Swap Byte Order Changes the byte order from Little Endian to Big Endian, or vice-versa. Some CD-ROM drives designed to work only with other types of computers (like DEC and Macintosh systems) report data by using the Little Endian byte order, while PCs use the Big Endian method. Normally, you should leave this box unchecked; check it *only* if the extraction process seems to work fine but the audio results are “garbage.”

Swap Channels Places the left channel of a CD's audio in the right channel of the Wave Display, and places the right channel of the audio in the Wave Display's left channel.

Spin Up Before Extraction Causes the CD-ROM drive to start spinning before Adobe Audition extracts the data. Some CD-ROM drives have better accuracy if they first read the CD after the drive is spinning. Selecting this option for other drives, however, doesn't provide any advantage.

Recording from CDs internally


If you have an older CD-ROM drive that doesn't support digital extraction, or if you have problems ripping a track into Adobe Audition, then you can record from a CD in real-time through the sound card on your computer. This method is called *internal recording*. Keep in mind that not all PC's have an analog cable from a CD drive, and not all computers react the same way when recording from CD internally. As a result, this method is never preferable to extracting from CD digitally.

Before you record from a CD internally, you should always preview the CD Audio input level to make sure that clipping won't occur.

To preview the CD Audio input level:

- 1 Open your favorite third-party CD player application (such as Windows Media Player).
- 2 Start playing the loudest part of the CD. Then, switch to Adobe Audition, and choose Options > Monitor Record Level.
- 3 Use the Level Meters in Adobe Audition to monitor the amplitude of the incoming signal. You want the input level to be as loud as possible without exceeding 0 dB. If the input level exceeds 0 dB, *clipping* occurs. (See "[Monitoring recording and playback levels](#)" on page 70.)
- 4 If you need to adjust the CD Audio input level, choose Options > Windows Recording Mixer to open the Windows Recording Control panel. Adjust the CD Audio input level as desired.
- 5 After you finish monitoring the input level, choose Options > Monitor Record Level.


To record from a CD internally:

- 1 In Edit View, create a new file.
- 2 Click the Record button .
- 3 Start the desired track in your CD player application.
- 4 When desired, stop recording in both Adobe Audition and the CD player application.

Setting the current-time indicator

The *current-time indicator* is a vertical, dotted line in the display window. You set the current-time indicator in order to start playback or recording at a specific point in a waveform.

When you work with multiple files in Edit View, you can use the Synchronize Cursor Across Windows command to retain the position of the current-time indicator between files. This command is useful if you switch between different versions of the same waveform during editing. In Multitrack View, you can use the Synchronize Clips With Edit View command to maintain the position of the current-time indicator when you switch between Multitrack View and Edit View.



Current-time indicator

To set the current-time indicator:

Do one of the following in the display window:

- Click exactly where you want to set the current time.
- Position the pointer over the triangle above or below the current-time indicator. (This triangle is the current-time indicator's *handle*.) Drag the handle to the desired position in the timeline.

After you set the current-time indicator, you can save it as a cue for later reference. For more information, see [“Working with cues” on page 82](#).

To synchronize the current-time indicator between waveforms:


In Edit View, choose Options > Synchronize Cursor Across Windows.

To synchronize the current-time indicator between Multitrack View and Edit View:

In Multitrack View, choose Options > Synchronize Clips With Edit View.

Monitoring time

Adobe Audition provides several features to help you monitor time during recording and playback. The *playback cursor*—a vertical, white line that appears in the display window—shows you the current time in the waveform. The Time window shows the current time in numerical format. The default display format is mm:ss:ddd (minutes:seconds:thousandths of a second), but you can easily change it. The display format is also used by the timeline along the bottom of the display window.




Features that help you monitor time

A. Playback cursor **B.** Timeline **C.** Time window

To display the Time window:

Do one of the following:

- Choose Window > Time. A check mark indicates that the window is visible.
- Click the Hide/Show Time Window button  in the View toolbar. (See [“Using toolbars” on page 22.](#))

If you don't like the default location of the Time window, you can reposition it or detach it so it floats above the main window. (See [“Using windows” on page 22.](#))

To change the time display format:

Choose View > Display Time Format, and choose the desired option:

- Decimal (mm:ss.ddd) displays time in minutes, seconds, and thousandths of a second.
- Compact Disc 75 fps displays time in the same format utilized by audio compact discs, where each second equals 75 frames.
- SMPTE 30 fps displays time in the SMPTE format, where each second equals 30 frames.
- SMPTE Drop (29.97 fps) displays time in the SMPTE drop-frame format, where each second equals 29.97 frames.
- SMPTE 29.97 fps displays time in the SMPTE non-drop-frame format, where each second equals 29.97 frames.
- SMPTE 25 fps (EBU) displays time using the standard European frame rate, where each second equals 25 frames.
- SMPTE 24 fps (Film) displays time in a format where each second equals 24 frames, suitable for film.
- Samples displays time numerically, using as a reference the actual number of samples that have passed since the beginning of the edited file.

- Bars and Beats displays time in a musical measures format of bars:beats:ticks. To adjust the settings, choose Edit Tempo. For more information, see [“Calculating the tempo of selected ranges” on page 216](#).
- Custom (X frames/sec) displays time in a custom format. To modify a custom format, choose Edit Custom Time Format, enter a number of frames per second for Custom Time Code Display, and click OK.

Using the transport controls

Just like many hardware-based audio recording and playback devices, Adobe Audition provides transport controls for playing, recording, stopping, pausing, fast forwarding, and rewinding waveforms and sessions.




Right-click the transport control buttons to set options for playing, recording, fast forwarding, and rewinding audio.



Transport controls

To show or hide the transport controls:

Do one of the following:

- Choose Window > Transport Controls. A check mark indicates that the controls are visible.
- Click the Hide/Show Transport Controls button  in the View toolbar. (See [“Using toolbars” on page 22](#).)

If you don't like the default location of the transport controls, you can reposition them or detach them so they float above the main window. (See [“Using windows” on page 22](#).)

Recording audio

You can record audio from a microphone or any signal you can plug into the Line In port of a sound card.



Note: You may need to adjust the input signal to obtain the optimum recording and signal-to-noise levels. (See [“Adjusting a sound card's levels” on page 72](#).)

By default, Adobe Audition displays waveforms in real time while recording. However, if the recorded audio is choppy, deselect Live Update During Recording in the General tab of the Settings dialog box. (See [“Setting Adobe Audition preferences” on page 43](#).)

Recording audio in Edit View

In Edit View, you can record audio into a new file or over existing audio. You can also disable the Record button so you don't start recording accidentally.

To record in Edit View:

- 1 Do one of the following:
 - Create a new file. (See [“Creating new audio files” on page 74.](#))
 - In an existing file, place the current-time indicator where you want to start recording. (See [“Setting the current-time indicator” on page 62.](#))
- 2 Click the Record button  to begin recording.
- 3 Click the Stop button  to stop recording.

To disable the Record button:

Right-click the Record button, and choose Disable Record Button. Repeat to reenable the button.


Using timed record mode

Use timed record mode to set start and stop times for recording. You can specify a maximum recording time and you can set a time for recording to start and stop automatically.

To enable or disable timed record mode:

Choose File > Timed Record Mode. Alternatively, right-click the Record button, and choose Timed Record Mode. A check mark indicates that timed record mode is enabled.


To set start and stop times for recording:

- 1 Enable timed record mode.
- 2 Click the Record button .
- 3 Specify the maximum recording time:
 - Select No Time Limit to record until you click the Stop button (or until disk space runs out).
 - Select Recording Length to record for the duration you specify in the Recording Length box.
- 4 Specify when to start recording:
 - Select Right Away to begin recording as soon as you click OK.
 - Select Time/Date to begin recording at a time you specify (for example, to have Adobe Audition capture a radio broadcast at a certain time). Enter the starting time and date in the appropriate text boxes, and set the desired time and date options.
- 5 Click OK.

Recording audio in Multitrack View

In Multitrack View, you can record audio on multiple tracks by overdubbing. When you overdub tracks, you can hear previously recorded tracks and play along with them to create sophisticated, layered compositions.

Each recording becomes a new audio clip on a track. If you are unsatisfied with a section of a recorded clip, you can select that section and punch in a new recording—leaving the remainder of the original clip intact. For particularly important or difficult sections, you can punch in multiple takes (different versions), and then select the take with the best performance.



A take created with the Punch In command

To record a new clip in a track:

- 1 In the track controls area for the track, click the In 1 button, select the desired input of your sound card, and then click OK.
- 2 Click the Record-enable button **R** for the track.
- 3 To simultaneously record on multiple tracks, repeat steps 1-2 for each track.
- 4 Position the current-time indicator at the desired starting point for recording, or select the range where you want to record the clip.
- 5 Click the Record button **R** to begin recording.
- 6 Click the Stop button **S** to stop recording.


To record in a loop:

- 1 Specify the input source, track, and starting point (or range) for recording, as described in the previous procedure.
- 2 Right-click the Record button **R**, and choose one of the following options:
 - Loop While Recording (View or Sel) to loop when the cursor reaches the end of the viewable range of track. If a range is selected, looping occurs when the cursor reaches the end of the range.
 - Loop While Recording (Entire or Sel) to loop when the cursor reaches the end of the track. If a range is selected, looping occurs when the cursor reaches the end of the range.
- 3 Click the Record button **R** to begin recording.
- 4 Click the Stop button **S** to stop recording.



If you use either of the Loop While Recording options for punching in audio, a new take is created with each loop.

To punch into a range of a clip:

- 1 In the track display, select the range of the clip.
- 2 Choose Edit > Punch In.
- 3 Position the current-time indicator a few seconds prior to the selected range.
- 4 In the Transport Controls window, click the Record button .
- 5 To punch in multiple takes, repeat step 4 for each take.

Note: You can't punch into a loop-enabled clip. For information about disabling loops, see ["Setting impermanent loop properties in Multitrack View" on page 218](#).

To select from multiple takes in a clip:

- 1 Select the clip.
- 2 Choose Edit > Take History, and then select the desired take.

To merge a selected take into a clip:

Choose Edit > Take History > Merge This Take (Destructive).

Note: Merging destructively adds a 30 millisecond crossfade at take edges.

To delete a selected take:

Choose Edit > Take History > Delete This Take.

Playing audio

Adobe Audition provides several ways to play audio, including using the transport controls to play the currently active file, using the Organizer window to preview files, and using the Windows Run command to start Adobe Audition and begin playing a file.


Playing audio by using the transport controls

The transport controls provide several options for playing the currently active file. For example, you can play just the visible section of a waveform, the duration from the current-time indicator to the end of the file, or the entire waveform. In addition, you can set preroll and postroll options to play a selection with just a bit of audio preceding or following it.



To start playback without using the transport controls, press the space bar. Press the space bar again to stop playback.


To play a range of audio:

Select the range you want to play, and click the Play button  in the Transport Controls window.

To play from the current-time indicator to the end of the current view:

Set the current-time indicator where you want playback to start, and click the Play button in the Transport Controls window.

To play from the current-time indicator to the end of the file:

Set the current-time indicator where you want playback to start, and click the Play To End button  in the Transport Controls window.

To play the visible portion of the file:


Right-click the Play button or Play To End button, and choose Play View. Then click the button again to start playback.

To play an entire file:

Right-click the Play button or Play To End button, and choose Play Entire File. Then click the button again to start playback.

To loop audio during playback:

Do one of the following:



- To play the currently-visible portion of the audio in a continuous loop, click the Play Looped button  in the Transport Controls window.
- To loop the entire waveform or session (or just the selected range), right-click the Play Looped button and choose Play Entire (or Selection). Then click the button again to start playback.

Note: By default, the display window scrolls in sync with playback that extends beyond the visible section of a waveform. In the General tab of the Settings dialog box, you can set options for auto-scrolling or you can disable this feature. (See [“Setting Adobe Audition preferences” on page 43.](#))

Using preroll and postroll during playback (Edit View only)

In Edit View, you can play back the audio just before and after a selected range. This audio is known as *preroll* and *postroll*. Playing preroll and postroll is useful for fine-tuning selections and listening to transitions without destroying a selection. By default, the duration of preroll and postroll is one second; however, you can adjust this duration to best meet your needs.



To play a selected range of audio with preroll and postroll:

- 1 In Edit View, right-click the Play button  or the Play To End button  in the Transport Controls window, and choose one of the following options: Play Preroll and Postroll, Play Preroll and Selection, Play Postroll, or Play Preroll, Postroll, and Selection
- 2 Click the button again to start playback.



You can also use keyboard shortcuts to play preroll and postroll. For information on specific keyboard shortcuts, [“Keys for playing audio” on page 264.](#)


To set a duration for preroll and postroll:

- 1 In Edit View, right-click the Play button  or the Play To End button  in the Transport Controls window.
- 2 Choose Preroll and Postroll Options.
- 3 In the Edit View–Play section of the Preroll and Postroll Options dialog box, specify a duration for preroll and postroll.
- 4 Click OK.

Previewing audio by using the Organizer window

The Files tab in the Organizer window provides several play options that make it easy to preview loops and other files. These options are particularly handy when you work in Multitrack View because they let you preview loops at the session tempo. For more information on using the Files tab in the Organizer window, see [“Organizing files” on page 30](#).

To preview a file:

- 1 Make sure that the advanced options—including the preview and sorting controls—appear in the Files tab of the Organizer window. If they don't, click the Advanced Options button  at the top of the Files tab.
- 2 Select the file you want to preview, and then click the Play button ▶. Click the Stop button to stop the preview. Use the volume slider to adjust the volume of the preview.

To enable auto-play:


Click the Auto-play button on the Files tab. Adobe Audition automatically previews files you select. To disable auto-play, click the Auto-play button again.

To preview a file at the session tempo (Multitrack View only):

In Multitrack View, select Follow Session on the Files tab, and then click the Play button or enable auto-play, and select a file.

Note: Only files that are loop-enabled can be previewed at the session tempo. Loop-enabled files are identified with a loop icon in the Files tab.

To enable continuous loop preview:

Click the Loop button  on the Files tab, and then click the Play button or enable auto-play, and select a file. To disable continuous loop preview, click the Loop button again.

Playing audio by using the Windows Run command

You can start Adobe Audition and begin playing a file by using the Windows Run command. Before using the command, make sure that Auto-Play On Command-Line Load in the General tab of the Settings dialog box is selected. (See [“Setting Adobe Audition preferences” on page 43](#).)

To play audio from the command line:

In the Windows Run dialog box, type the following text, and click OK:

"[drive]:\Program Files\Adobe\Audition 1.5\Audition.exe" "[path to file]"

For example, type the following to play the TalkBackVerb loop file:

"c:\Program Files\Adobe\Audition 1.5\Audition.exe" "c:\Program Files\Adobe\Audition 1.5\Audition Theme\TalkBackVerb.cel"

Stopping, pausing, and adjusting the playback cursor

The transport controls provide buttons for stop recording and playback, pausing recording and playback, and adjusting the playback cursor.



To stop playback without using the transport controls, press the spacebar. Press the spacebar again to start playback.

To stop playing or recording audio:

Click the Stop button ■ in the Transport Controls window.

To pause playing or recording audio:

Click the Pause button || in the Transport Controls window. Click the Pause button again to resume playback or recording.

To adjust the playback cursor:

Click one of the following buttons in the Transport Controls window:

- The Go to Beginning button ◀◀ places the playback cursor at the beginning of the waveform or session.
- The Rewind button ◀◀ shuttles the playback cursor backward in time. This function supports scrubbing, meaning that on some sound cards, the audio file plays back at a lower volume as the playback cursor shuttles over the waveform or session.



Right-click the Rewind button to set the rate at which the cursor moves.

- The Fast Forward button ▶▶ shuttles the playback cursor forward in time. This function supports scrubbing, meaning that on some sound cards, the audio file plays back at a lower volume as the playback cursor shuttles over the waveform or session.



Right-click the Fast Forward button to set the rate at which the cursor moves.

- The Go to End button ▶▶ places the playback cursor at the end of a waveform (in Edit View) or at the end of the list clip in a session (in Multitrack View).

Monitoring recording and playback levels

Adobe Audition provides the Level Meters to help you monitor the amplitude of the signal during recording and playback. If the amplitude is too high, *clipping* occurs and results in distortion; if the amplitude is too low, the sound quality is reduced.

If you find that the signal is too high or low during recording and playback, you can adjust the input and output levels of your sound card.


Using the Level Meters

The Level Meters represent the incoming signal in *dBFS* (decibels below full scale), where a level of 0 dB is the maximum amplitude possible before clipping occurs. Yellow peak indicators remain for 1.5 seconds to allow for reading of the peak amplitude. If clipping does occur, the clip indicator to the right of the meter lights up and stays on until you clear it. When stereo audio is displayed, the top meter represents the left channel, and the bottom represents the right.

You can customize the Level Meters in a variety of ways, such as changing the decibel range, showing valley (minimum amplitude) indicators, and changing the mode of the peak indicators.



Right-click the Level Meters to set metering options.




The Level Meters

A. Left channel **B.** Right channel **C.** Peak indicators **D.** Clip indicators

To show or hide the Level Meters:

Do one of the following:

- Choose Window > Level Meters. A check mark indicates that the Level Meters are visible.
- Click the Hide/Show Level Meters button  in the View toolbar. (See [“Using toolbars” on page 22.](#))

If you don't like the default location of the Level Meters, you can reposition them or detach them so they float above the main window. (See [“Using windows” on page 22.](#))

To start or stop monitoring the levels of an input source:

Choose Options > Monitor Record Level, or double-click the Level Meters.

To disable or enable the Level Meters during recording or playback:

Choose Options > Show Levels On Play And Record. Disabling the Level Meters improves performance on lower specification computers.

To clear a clip indicator:

Click the clip indicator, or right-click the Level Meters and choose Clear Clip Indicator.

Note: The clip indicators always light up if clipping occurs, but if Adjust For DC is enabled, the indicators light up if audio has a DC offset.

To adjust for DC offset:

Right-click the Level Meters, and choose Adjust For DC.

Many sound cards record audio with a slight DC offset, meaning that the center of the waveform being recorded is a little above or below the center of the waveform display. This offset can dramatically throw off the level meters since the offset amount could be interpreted as a constant sound at that volume. You should have this option enabled when recording.

To show or hide valley indicators:

Right-click the Level Meters, and choose Show Valleys.

If the valley indicators are close to the peak indicators, the dynamic range (the difference between the quietest and loudest sounds) is low. If they're spread far apart, the dynamic range is high.

To change the decibel range of the Level Meters:

Right-click the Level Meters, and choose a Range option.

To change the mode of peak indicators:

Right-click the Level Meters, and choose one of the following options:

- Dynamic Peaks causes the yellow peak level indicators to reset to a new peak level after 1.5 seconds, letting you easily see the peak amplitude “right now.” As the audio gets quieter, the peak indicators start backing off.
- Static Peaks keeps the peak levels from being reset, letting you retain the maximum amplitude of the signal since monitoring, playing, or recording began. The peak can still be reset manually at any time by clearing the clip indicators (that is, by clicking the clip indicator at the right).



Select Static Peaks as a great way to find out how loud a song will get before you record it. Just start monitoring levels and then play the song. After the song ends, the peak indicators show the volume of the loudest part.

Adjusting a sound card's levels

Adobe Audition doesn't directly control a sound card's record levels (input gain) and playback levels (output volume). Instead, you can adjust these levels with the mixer application that comes with the sound card or with the mixer built into Windows. You may need to adjust levels if recordings are too quiet (causing unwanted noise), too loud (leading to clipped, distorted sound), or not audible when played in Adobe Audition.



To get the best sounding results, you should record audio as loud as possible without clipping. Try to keep the loudest point somewhere between -2 dB and 0 dB when setting the recording levels.

To adjust a sound card's record and playback levels by using Windows:

1 Open the Windows Volume Control program.

You can usually access this program in the Programs > Accessories > Entertainment (or Multimedia) menu of the Windows Start menu. On many systems, you can also double-click the speaker icon in the system tray to access the Volume Control program, which resembles a small mixing board with vertical sliders.

2 To adjust the sound card's playback (output) level, turn up the sliders on the Windows mixer to the desired volume. Make sure that Mute underneath both sliders isn't selected.

3 To adjust the sound card's record (input) level, choose Options > Properties in Volume Control. Select Recording and click OK. Be sure that the input source you want to use is selected, and adjust other sliders on the Windows mixer as needed.



To quickly access the Record section of the Windows mixer, choose Options > Windows Recording Mixer in Adobe Audition.


Editing Audio

About editing audio

When you open an audio file in Edit View, you see the *waveform display*, a visual representation of the sound wave, or *waveform*. If you open a stereo file, the waveform for the left channel appears at the top and the waveform for the right channel appears at the bottom. If you open a mono file, the waveform utilizes the total height of the waveform display. The peaks and valleys in the waveform represent positive and negative air pressure. Quiet audio has both lower peaks and lower valleys than loud audio.




For background information on working with digital audio, see [“Sound fundamentals” on page 267](#).



Stereo waveform in Edit View

Many editing tasks require that you select a precise range of a waveform. When selecting a range, you'll probably want to zoom in to view the waveform in more detail. (See [“Zooming” on page 25](#).) Adobe Audition provides a variety of ways to select audio data precisely, such as adjusting selections to zero-crossings, finding beats, and using snapping. (See [“Selecting audio data” on page 76](#).)

As you edit a waveform, keep in mind that you can undo your changes until you save the file. (See [“Undoing and redoing changes” on page 29](#).)



Creating new audio files

The File > New command lets you create an empty audio file. Doing so is useful when you want to paste audio into an empty file before you edit it or when you want to record audio into a new file.



You can quickly create a new file from a selection by choosing Edit > Copy to New. (See [“Copying audio data” on page 80.](#))

To create a new audio file:

1 Choose File > New. Alternatively, click the New File button  in the toolbar.

2 Select a sample rate in the list, or type a custom sample rate in the text box.

The sample rate determines how many frequencies can be encoded in the audio signal. (Higher sampling rates mean a wider bandwidth.) For more information, see [“About sample rates” on page 93.](#)

3 Select a number of channels:

- Mono creates a waveform with just one channel of audio information. This setting is good for voice-only recordings.
- Stereo creates a waveform with separate right and left channels. This setting is usually best for music recordings. Because stereo waveforms contain twice as much data as mono waveforms, they consume twice as much storage space.

4 Select a resolution, and click OK:

- 8-bit creates a waveform where quality is not much of a concern, but small file size is. 8-bit waveforms are usually fine for telephony applications or for embedded sounds in Web pages. Although they tend to be noisier than their 16-bit counterparts, they're half the size.
- 16-bit produces a CD-quality waveform. This setting is suitable for most broadcast and music recordings.
- 32-bit creates a waveform that supports the most precise audio processing, and 32-bit is the recommended resolution for editing files in Adobe Audition. After you edit a file, you can downsample it to 16- or 8-bit for output and achieve better results than if you edit an 8- or a 16-bit file. (See [“Changing the bit depth” on page 96.](#))


Note: Older sound cards might not be able to play 32-bit files properly. To check the capabilities of your sound card, choose Options > Device Properties. If your sound card doesn't support 32-bit files, you can limit playback to 16-bits while retaining the 32-bit depth internally. (See [“Setting properties for audio output devices” on page 38.](#))

Viewing waveforms

The waveform display in Edit View shows you a visual representation of a waveform. By default, this representation shows the amplitude of a waveform over time. However, you can view the frequency of a waveform over time by switching to Spectral View. You can also control the scale with which Adobe Audition measures the amplitude or frequency of waveforms.

Switching between Waveform View and Spectral View


The waveform display offers two ways to represent audio data: Waveform View and Spectral View.



Waveform View and Spectral View

- Waveform View (the default) displays a waveform as a series of positive and negative peaks. The x-axis (horizontal ruler) represents time, and the y-axis (vertical ruler) measures spikes, or increased amplitude, in a waveform."
- Spectral View displays a waveform by its frequency components, where the x-axis represents time and the y-axis (vertical ruler) measures frequency. This view lets you analyze audio data to see which frequencies are most prevalent. The greater a signal's amplitude component within a specific frequency range, the brighter the displayed color. Colors range from dark blue (meaning that the frequencies are very low in amplitude) to bright yellow (meaning that the frequencies are high in amplitude).

To switch between Waveform View and Spectral View:

Choose View > Waveform View or View > Spectral View. Alternatively, click the Toggle Between Waveform And Spectral Views button  in the toolbar.

Adobe Audition lets you customize certain features of Waveform View and Spectral View. For example, you can show or hide grid lines in Waveform View and change the resolution in Spectral View. For more information, see ["Display options" on page 48](#).

Changing the vertical scale

In Waveform View, the vertical ruler shows the decibel value of the audio data. However, you can change the scale of the ruler to Sample Values, Normalized Values, or Percentage.

Note: In Spectral View, the vertical scale is always in hertz (Hz).

To change the scale of the vertical ruler:

Choose View > Vertical Scale Format, and choose the desired scale:

- Sample Values indicates amplitude as the data's exact sample value of the data.
- Normalized Values indicates amplitude on a normalized scale value that ranges from -1 to 1.
- Percentage indicates amplitude on a percentage scale value that ranges from -100% to 100%.
- Decibels indicates amplitude using a decibel scale value that ranges from -Infinity to Zero.



Double-click the vertical ruler to cycle through the scales.

Selecting audio data


To edit a waveform, you must first select the audio data that you want to modify. Adobe Audition provides several methods for making and adjusting selections.



Using cues ranges can save you time when making selections. (See [“Working with cues” on page 82.](#))

Selecting with the mouse

You can select a range of audio data by dragging in the waveform display. When precision is important, you may want to zoom in to view the waveform in more detail. (See [“Zooming” on page 25.](#))



Dragging to select a range

To select a range of a waveform:

Drag to select the desired range of the waveform.

To extend or shorten a selection:

Shift-click the end of the selection that you wish to modify, and drag to extend or shorten it.

Note: If you prefer, you can right-click to extend or shorten a selection. To enable this feature, select *Extend Selection* in the *General* tab of the *Settings* dialog box. (See [“Setting Adobe Audition preferences” on page 43.](#))

To select a range in only one channel:

Do one of the following:

- Drag near the top of the left (upper) channel. The cursor displays an *L* icon to indicate the left channel.
- Drag near the bottom of the right (lower) channel. The cursor displays an *R* icon to indicate the right channel.

To select the visible range of a waveform:

Double-click in the waveform display.

Selecting all of a waveform


The *Select Entire Wave* command lets you select all of the audio data in a waveform. You can use the mouse to do this as well.

To select all of a waveform:

Choose *Edit > Select Entire Wave*, or triple-click in the waveform display.


Selecting audio frequencies in Spectral View

When working in Spectral View, you can use the Marquee Selection tool to select audio data within specific frequencies. This method allows for band-limited editing and processing, as well as greater flexibility in audio restoration work. For example, if you detect an audio anomaly or error, you can select and edit just the affected frequencies, resulting in superior results and faster processing.



Making a marquee selection in Spectral View

To make a marquee selection:

- 1 In Spectral View, click the Marquee Selection button  in the toolbar. If this button isn't visible, choose View > Toolbars > Spectral Selection.
- 2 Drag in the waveform display to select the desired audio data.



When making a marquee selection in a stereo waveform, the selection is applied to both channels. To select audio data in just one channel, choose Edit > Edit Channel, and then choose Edit Left Channel or Edit Right Channel.

To move a marquee selection:

Position the cursor in the selection, and drag it to the desired location.

To resize a marquee selection:


Position the cursor on the corner or edge of the selection, and drag it to the desired size.

Adjusting selections to zero-crossing points

For many editing tasks, such as deleting or inserting audio in the middle of a waveform, the best places to make selections are the points where the amplitude is zero (called *zero-crossings*). Selecting the zero-crossing points reduces the chance that an edit will create an audible pop or click. You can easily adjust a selection to the closest zero-crossing points by using a Zero Crossing command.

To adjust a selection to zero-crossing points:

Choose Edit > Zero Crossing, and choose one of the following commands:

- Adjust Selection Inward adjusts both range boundaries inward to the next zero-crossing point. Alternatively, click the Zero Crossing button  in the toolbar.
- Adjust Selection Outward adjusts both range boundaries outward to the next zero-crossing point.
- Adjust Left Side To Left adjusts the left range boundary leftward to the next zero-crossing point.

- Adjust Left Side To Right adjusts the left range boundary rightward to the next zero crossing point.
- Adjust Right Side To Left adjusts the right range boundary leftward to the next zero crossing point.
- Adjust Right Side to Right adjusts the right range boundary rightward to the next zero crossing point.

Finding beats

For some editing tasks, such as constructing drum loops and similar musical phrases, you need to select audio between beats. You can usually pick out where the beats are by looking for the peaks in a waveform. You can also use a Find Beats command.



Once you find beats, you can save them as Beat Cues, making it easy to locate the beats again. (See [“Working with cues” on page 82.](#))

To find the beginning of a beat:

- 1 Click in the waveform display to the left of the first beat you want to find.
- 2 Choose Edit > Find Beats > Find Next Beat (Left). The cursor moves to the beginning of the next beat.

To select audio between beats:

- 1 Find the beginning of a beat.
- 2 Choose Edit > Find Beats > Find Next Beat (Right) to select from the current cursor position to the next beat.
- 3 If you want to select more than one beat, choose Edit > Find Beats > Find Next Beat (Right) again. Each time you choose this command, Adobe Audition adds the next beat to the selection.



When you select audio to construct a loop, click the Play Looped button in the transport controls to preview the loop. After any necessary tweaking, you can then save, paste, or add the loop to the Cue List.

To adjust the settings that Adobe Audition uses to find beats:

Choose Edit > Find Beats > Beat Settings. Enter new values for Decibel Rise and Rise Time, and click OK.



For finding beats with material that has fast transient attacks, such as drums, specify a quick Rise Time and a high Decibel Rise so as not to cut off the beginning of the attack. For material with softer attacks, such as bass, the Rise Time can be slightly slower relative to Decibel Rise.

Snapping

Snapping causes selection boundaries, as well as the current-time indicator, to move to items such as cues, ruler ticks, zero-crossing points, and frames. Enabling snapping helps you make accurate selections; however, if you prefer, you can disable snapping for specific items.

To enable or disable snapping:

Choose Edit > Snapping, and choose any of the following commands. A check mark indicates that a command is enabled:

- Snap To Cues allows the cursor to snap to a cue point. For more information on defining cues, see [“Working with cues” on page 82](#).
- Snap To Ruler (Coarse) allows the cursor to snap only to the major numeric divisions (decimal, SMPTE, samples, and so on) in the timeline.

Note: You can enable only one Snap To Ruler command at a time.

- Snap To Ruler (Fine) allows the cursor to snap to each of the subdivisions (decimal, SMPTE, samples, and so on) within the timeline. Zooming in (by right-clicking as you drag across the timeline) breaks the display down into more accurate subdivisions, letting you place the cursor more accurately within the timeline.
- Snap To Zero Crossings allows the cursor to snap to the nearest place where the waveform crosses the center line (in other words, the zero amplitude point).
- Snap To Frames (Always) allows the cursor to snap to a frame boundary, as long as the time format is measured in frames (such as Compact Disc and SMPTE). This command is especially handy for working on audio for CD.






You can access snapping commands by right-clicking the timeline.

Specifying which channel of a stereo waveform to edit

By default, Adobe Audition applies selections and edits to both channels of a stereo waveform. However, you can easily select and edit just the left or right channel of a stereo waveform.

To specify which channel you want to edit:

Do either of the following:

- Choose Edit > Edit Channel, and choose which channel you want to edit.
- Click the Edit Left Channel button , the Edit Right Channel button , or the Edit Both Channels button  in the View toolbar. (See [“Using toolbars” on page 22](#).)

Copying, cutting, pasting, and deleting

Adobe Audition supports all of the basic editing functions, as well as several commands designed specifically for audio editing.

Choosing a clipboard

Adobe Audition gives you access to five internal clipboards for temporary data storage. Each works similarly to the Windows clipboard, except that they can handle more data at a faster rate.

To choose a clipboard:

Choose Edit > Set Current Clipboard, and choose a clipboard.




Choose the Windows clipboard if you want to copy audio data to other Windows applications.

Copying audio data

The Copy command lets you copy audio data to the active clipboard. The Copy To New command lets you copy and paste the data to a new file in one step.


To copy audio data:

- 1 In the waveform display, select the audio data you want to copy. Or, to copy the entire waveform, deselect all audio data.
- 2 Choose Edit > Copy or Edit > Copy To New. Alternatively, click the Copy button  in the toolbar.

Cutting audio data

The Cut command lets you remove audio data from the current waveform and copy it to the active clipboard.

To cut audio data:

- 1 Select the audio data you want to cut. Or, to cut the entire waveform, deselect all audio data.
- 2 Choose Edit > Cut. Alternatively, click the Cut button  in the toolbar.

Pasting audio data


The Paste command lets you paste audio data from the active clipboard to the current waveform. If the format of the data on the clipboard differs from the format of the file it's being pasted into, Adobe Audition automatically converts the format before pasting the data.

The Paste To New command lets you create a new file and insert audio data from the active clipboard. The new file automatically inherits the properties (sample rate, sample frequency, and so on) from the original clipboard material.



The Highlight After Paste option in the General tab of the Settings dialog box determines whether or not data is highlighted after you paste it into a file.

To paste audio data into the current file:

- 1 In the waveform display, place the cursor where you want to insert the audio data or select the audio data you want to replace.
- 2 Choose Edit > Paste. Alternatively, click the Paste button  in the toolbar.

To paste audio data into a new file:

Choose Edit > Paste To New.

Mixing audio data when pasting

The Mix Paste command lets you mix audio data from the clipboard or a file with the current waveform. If the format of the data on the clipboard differs from the format of the file it's being pasted into, Adobe Audition automatically converts the format before pasting the data.



The Mix Paste command provides a quick alternative to using the more powerful and flexible multitrack functions in Adobe Audition.

To mix audio data with the current waveform:

1 In the waveform display, place the cursor where you want to start mixing the audio data. Alternately, select the audio data you want to replace.

2 Choose Edit > Mix Paste. Alternatively, click the Mix Paste button  in the toolbar.

3 Set the following options as desired, and click OK.

Volume Adjusts the sound level of the left and right channels before pasting. Move the volume sliders, or enter a percentage in the text boxes to the right of them.



Paste in single channels (either left or right) by adjusting the level of the opposite channel to zero.

Invert Turns that channel of the waveform upside-down. (Any samples above the center line are placed below it, and those below the center line are placed above it.)



This option is handy when you want to take the difference between two samples (or subtract one signal from another).

Lock Left/Right Locks the volume sliders so that they move together.

Insert Inserts audio at the current location or selection, replacing any selected data. If no data is selected, Adobe Audition inserts audio at the cursor location, moving any existing data to the end of the inserted material.

Overlap Mixes audio at the selected volume level with the current waveform. If the audio is longer than the current waveform, the current waveform is lengthened to accommodate the pasted audio.

Replace Overdubs the audio beginning at the cursor location, and replaces the existing material thereafter for the duration of audio. For example, pasting 5 seconds of material replaces the first 5 seconds after the cursor.

Modulate Modulates the audio with the current waveform for an interesting effect. The result is similar to overlapping, except that the values of the two waveforms are multiplied by each other, sample by sample, instead of added.



You can create fantastic combo effects by selecting part of a wave and using the Mix Paste command with Modulate selected. The selection is modulated with the audio signal on the clipboard.

Crossfade Applies a fade to the beginning and end of the pasted audio. Enter a value to specify how many milliseconds of the audio are faded.



Use this option for smoother transitions to and from pasted audio.

From Clipboard [number] Pastes audio data from the active internal clipboard.

From Windows Clipboard Pastes audio data from the Windows clipboard. If the Windows clipboard contains no audio data, this option is disabled.

From File Pastes audio data from a file. Click Select File to browse for the file.


Loop Paste Pastes audio data the specified number of times. If the audio is longer than the current selection, the current selection is automatically lengthened accordingly.

Deleting audio data


Adobe Audition provides two methods for deleting audio: The Delete Selection command lets you remove a range from a waveform, whereas the Trim command lets you remove unwanted audio from both sides of the selected audio.

Note: Deleted data doesn't go to the clipboard and can be retrieved only by choosing *Edit > Undo* or *File > Revert To Saved*, but only if you haven't saved the file since deleting the data.

To delete audio data:

- 1 In the waveform display, select the audio data you want to delete.
- 2 Choose *Edit > Delete Selection*. Alternatively, click the Delete button  in the toolbar.

To trim audio data:

- 1 In the waveform display, select the audio data you want to keep.
- 2 Choose *Edit > Trim*. Alternatively, click the Trim button  in the toolbar.


Working with cues

Cues are locations in a waveform that you define. Using cues makes it easy to navigate within a waveform in order to make a selection, perform edits, or play back audio.

About cues

In Adobe Audition, a cue can be either a *point* or a *range*. A point refers to an exact position within a waveform (for instance, 1:08.566 from the start of the wave). A range has both a start time and an end time (for example, all of the waveform from 1:08.566 to 3:07.379). If a cue is a range, you can drag its beginning and end points to different times.

Cues have triangular handles that appear at the top and bottom of the waveform display. You use cue handles to select and adjust cues. You can also right-click a cue handle to view commands for working with cues.



Examples of cues


A. Cue handle **B.** Cue point **C.** Cue range **D.** Nonsplit cue range

Note: To preserve cues when you save a file, make sure that you select *Save Extra Non-Audio Information*.

Defining and selecting cues

You use the Cue List window to define and select cues. You can also define cues by using context commands and keyboard shortcuts.


To display the Cue List:

Choose Window > Cue List. Alternatively, click the Hide/Show Cue List button  in the toolbar.

To define a cue:

1 Do one of the following:

- Place the cursor exactly where you want the cue point to be in the waveform display.
- Select the audio data you want to define as a cue range in the waveform display.

2 Click Add in the Cue List window. Alternatively, click the Add To Cue List button  in the toolbar.

To select cues:

Do one of the following:

- Click a cue in the cue list.
- Double-click a cue handle in the waveform display.
- To select adjacent (*contiguous*) cues, click the first cue you want to select in the cue list, and then Shift-click the last.
- To select nonadjacent (*noncontiguous*) cues, Ctrl-click them in the cue list.

Playing cues

The Auto Play feature in the Cue List window causes Adobe Audition to automatically play cues when you select them.

To enable or disable the Auto Play feature for cues:

Click the Auto Play button in the Cue List window.

Choosing a cue type

Adobe Audition provides four cue types. All four can be ranges as well as points, although it doesn't really make sense for index cues to be ranges. Consider the following when choosing a cue type:

- Basic Cues lets you mark important sections of a waveform for later reference (for example, to remind yourself of an editing point). Basic cues are also useful for specifying stop and start positions for the play list.
- Beat Cues is just like Basic Cues, but you use it to mark musical beats. Beat cues are a very powerful feature because an audio file saved with them allows the beat mapping loop method to be very accurate. For more information on creating and using loops, see ["About loops" on page 214](#).
- Track Cues lets you indicate a split in tracks for an audio compact disc. Use these cues only for burning CDs. (See ["Inserting tracks" on page 259](#).)
- Index Cues lets you set markers within a CD track. (Some CD players offer controls for cueing indexes.) Also, the time between the track cue that begins a track and the first index cue in that track shows up on the player as "negative time."

To change the cue type:

- 1 Select a cue.
- 2 Click Edit Cue Info in the Cue List window.
- 3 Choose a cue type from the Type menu.

Alternatively, right-click the cue handle, and choose a cue type from the context menu.

Naming cues

After you create a cue, you can rename it and add descriptive information.

To rename a cue and add a description:

- 1 Select a cue.
- 2 Click Edit Cue Info in the Cue List window.
- 3 Do either or both of the following:
 - Enter a new name in the Label text box.
 - Enter a description in the Desc text box.

Adjusting cues

You can easily adjust the position of cues, as well as the duration of range cues.

To reposition a cue:

Do one of the following:

- For point cues, drag the cue handle to a new location in the waveform display.
- For range cues, drag the red start handle to a new location in the waveform display.
- Select the cue, and click Edit Cue Info in the Cue List window. Enter a new value in the Begin text box.

To change the duration of a range cue:

- Drag the blue end handle to a new location in the waveform display.
- Select the cue, and click Edit Cue Info in the Cue List window. Enter a new value in the End or Length text box.

Merging, converting, and deleting cues

Adobe Audition lets you merge cues, and it also lets you convert point cues to range cues, and vice versa. If you find that you don't need a cue, you can delete it.

To merge cues:

- 1 Select the cues you want to merge. You can select only two cue ranges to merge, but you can select any number of cue points.
- 2 Click Merge in the Cue List window.

Note: The new merged cue inherits its name from the first cue. You lose any information in the Label and Desc text boxes for the subsequent merged cue.

To convert a point cue to a range cue:

Right-click the cue handle, and choose Make Range. The cue handle splits into two handles.

To convert a range cue to a point cue:

Right-click a cue handle, and choose Make Point. The two parts of the range cue handle merge into a single handle, with the start time of the range becoming the time for the point cue.

To delete cues:

- 1 Select one or more cues.
- 2 Click Del in the Cue List window. Alternatively, right-click the cue handle, and choose Delete.

Batch processing cues

You can use the Batch feature in the Cue List window to add silence between cues and save the audio between cues to new files.

To batch process cues:

- 1 Select one or more cues. At least one of the cues you select must be a range.
- 2 Click Batch in the Cue List window.
- 3 Set the following options as desired, and click OK:

Set Amount Of Silence Adds silence between cue points in the waveform. Enter the desired values (measured in seconds) in the Add Silence Before and Add Silence After text boxes.

Save To Files Saves the audio between cue points to new files.

Filename Prefix Specifies a prefix for the new files. Adobe Audition automatically adds numbers after the prefix (phrase02, phrase03, and so on) as well as the correct extension based upon the output format you choose.

Destination Folder Specifies the folder where Adobe Audition places the new files. Click Browse to specify a different folder.

Output Format Specifies the desired output format for the new files. If the specified format has options, the Options button is enabled. Click this button to select options.

Setting cues automatically

The Auto-Cue feature lets you locate phrases or beats and automatically add them to the cue list. You can also use this feature to remove silence from the beginning and end of a file.

To set cues automatically:

- 1 Select the general range in which you want to find phrases or beats.
- 2 Choose Edit > Auto-Cue, and choose one of the following commands:
 - Adjust Selection To Phrase selects a phrase within the selected range by adjusting the highlight inward, ignoring any silence before and after the audio. Nothing is added to the cue list.
 - Find Phrases And Mark scans the selected range, marking nonsilent ranges as basic cues in the cue list.
 - Find Beats And Mark scans the selected range, marking beats as beat cues in the cue list.

To customize Auto-Cue settings:

1 Choose Edit > Auto-Cue > Auto-Cue Settings.

2 Adjust the following options, and click OK:

- Audio Will Be Considered “Silence” When specifies parameters for finding silence. In the Signal Is Below text box, enter the amplitude value (in decibels) you want Adobe Audition to consider as the maximum level for silence. In the For More Than text box, enter the duration (in milliseconds) of this maximum amplitude value.



For very quiet, high-quality audio, enter a lower amplitude value (such as –60 dB). For noisier audio, the value might be much higher (such as –30 dB). Enter a longer duration to keep groups of words together, for example.

- Audio Will Be Considered As Valid When specifies parameters for determining if audio is valid. In the Signal Is Above text box, enter the amplitude value (in decibels) you want Adobe Audition to consider as the minimum level for audio. In the For More Than text box, enter the duration (in milliseconds) of this minimum amplitude.



Enter a longer duration to ignore short periods of undesired audio (like clicks, static, or other noise). However, if the value is too high (above 200 milliseconds), short words may be skipped.

- Find Levels scans the waveform (or a selected range) to have Adobe Audition automatically determine a good starting point for signal levels. Suggested values appear in the appropriate text boxes.



If these values don't do the job—for example, words or phrases get chopped off—lower the signal level values. Increase the signal level values if not enough silence is removed.

To trim silence from the beginning and ending of a file:

Choose Edit > Auto-Cue > Trim Digital Silence.

If you select the middle of a waveform, this command functions like the normal Trim command, trimming out everything else, in addition to any digital silence in the highlighted range at the endpoints.

Creating play lists

A *play list* is an arrangement of cue ranges that you can play back in any order and loop a specified number of times. The advantage of using a play list is that you can try different versions of an arrangement before you commit to the edits. You create play lists in the Play List window.

To display the Play List window:

Choose Window > Play List. Alternatively, click the Hide/Show Play List button  in the toolbar.

To create a play list:

1 If the Cue List window isn't visible, click Show Cue List in the Play List window.

2 In the Cue List window, select the cue ranges you want to add to the play list. (See [“Defining and selecting cues” on page 83.](#))

3 Click Insert Cues in the Play List window. The selection is inserted either before the currently selected item or at the end if nothing is selected.

To play items in a play list:

Do one of the following:

- To play the entire play list, select the first item in the list, and click Play in the Play List window.
- To play part of the list, select the first item you want to play, and click Play in the Play List window.
- To play a specific item in the list, select that item, and click Autocue in the Play List window.

To change the order of items in a play list:

- 1 Select the item you want to move.
- 2 Click Move Up or Move Down.

To set up looping for an item in a play list:

Select the item, and enter a number in the Loop text box. Each item in the play list can loop a different number of times.

To delete items from a play list:

Select the items you want to delete, and click Remove in the Play List window.

Creating and deleting silence

Adobe Audition provides several ways to create silence in and delete silence from a waveform. Creating silence is useful for inserting pauses and removing nonessential noise from an audio file. Removing silence is useful for cleaning up voice prompts and speeding up narratives without affecting the foreground audio.

Creating silence

Adobe Audition provides two ways to create silence in a waveform: by muting part of the existing waveform or by inserting a new duration of silence.

To mute existing audio data:

- 1 Select the desired range of audio data.
- 2 Choose Effects > Silence.

Unlike deleting or cutting a selection, which splices the surrounding material together, applying the Silence effect leaves the duration of the selection intact, and simply zeros the amplitude within it.

To insert a new duration of silence:

- 1 Place the cursor where you want to insert the silence. Or, if you want to replace part of the existing waveform, select the desired range of audio data.
- 2 Choose Generate > Silence.
- 3 Enter the number of seconds of silence you want to generate. Use decimals to enter partial seconds. For example, enter .3 to generate three-tenths of a second of silence.
- 4 Click OK. Any audio to the right of the cursor is pushed out in time, thereby lengthening the waveform's duration.

Deleting silence

The Delete Silence command detects and removes silence between words or other audio. It's ideal for cleaning up voice prompts and speeding up narratives without affecting the foreground audio.

To delete silence:

- 1 If you want to delete silence from part of a waveform, select the desired range of audio data. If you don't select a range, Adobe Audition deletes silence from the entire waveform.
- 2 Choose Edit > Delete Silence.
- 3 Set the following options as desired, and click OK:

"Silence" Is Defined As Determines what Adobe Audition considers silence. In the Signal Is Below text box, enter the amplitude value (in decibels) you want Adobe Audition to consider as the maximum level for silence. In the For More Than text box, enter the duration (in milliseconds) of this maximum amplitude value.



For very quiet, low-noise-floor audio, enter a lower amplitude value (such as -60 dB). For noisier audio, you might enter a higher value (such as -30 dB). Enter a longer duration to keep groups of words together, for example.

"Audio" Is Defined As Determines what Adobe Audition considers audio. In the Signal Is Above text box, enter the amplitude value (in decibels) you want Adobe Audition to consider as the minimum level for audio. In the For More Than text box, enter the duration (in milliseconds) of this minimum amplitude value.



Enter a higher duration to ignore short periods of undesired audio (like clicks, static, or other noise). However, if the value is too high (above 200 milliseconds), short words might be skipped.

Find Levels Scans the waveform (or selected range) to have Adobe Audition automatically determine a good starting point for signal levels. Suggested values appear in the appropriate text boxes.



If these values don't do the job—for example, words or phrases are chopped off—lower the signal level values. Increase the signal level values if not enough silence is removed.

Mark Deletions In Cue List Adds each location where silence is removed to the cue list.

Limit Continuous Silence To Specifies the minimum amount of silence (in milliseconds) to keep at all times. Silent ranges shorter than this length aren't removed; silent ranges greater than this length are shortened so that exactly the specified amount of silence remains. Set this value to zero to remove as much silence as possible.



When shortening speech segments, use a setting of 150 milliseconds or so to leave a more realistic, natural sounding pause. Higher values can lead to an artificial sounding pause.

Scan For Silence Now Previews the silence to be removed. This option reports how much silence will be removed and how many sections of silence were found. This option doesn't actually remove silence, but it gives you an idea of what to expect with the current settings when you actually choose the Delete Silence command.



If you have an audio presentation that consists of many cuts separated by silence (such as a reel of several jingles), choose Edit > Delete Silence to make sure that the duration of silence between each cut is the same. For example, if the difference between cuts 1 and 2 is 3.2 seconds, the difference between cuts 2 and 3 is 4.1 seconds, and the difference between cuts 3 and 4 is 3.7 seconds, you can use Delete Silence to make the silence duration between all four cuts exactly 3 seconds.

Inverting and reversing audio

The Invert effect simply inverts the waveform's samples, making all positive offsets negative and all negative offsets positive. Inverting is useful for lining up amplitude curves when creating loops or pasting. By inverting one channel of a stereo recording, you can also correct out-of-phase channels or create interesting phasing effects. For more information on phase, see the Glossary.

The Reverse effect reverses the order of a waveform's samples so that they play backwards. Reversing is useful for creating special effects.

To invert a waveform:

- 1 If you want to invert part of a waveform, select the desired range. Otherwise, deselect all audio data to invert the entire waveform.
- 2 Choose Effects > Invert.

To reverse a waveform:

- 1 If you want to reverse part of the waveform, select the desired range. Otherwise, deselect all audio data to reverse the entire waveform.
- 2 Choose Effects > Reverse.

Generating audio

Adobe Audition provides several commands that let you generate new audio data. These commands are different from effects in that they insert new sounds into a waveform rather than alter existing sounds.

Generating DTMF signals

Dual Tone Multi-Frequency (DTMF) signals (also known as touch tones) are used for dialing telephone numbers over phone lines that are capable of responding to touch tone signals. These signals are recommended internationally by the International Telegraph and Telephone Consultative Committee as the signals for push button telephones.

Keep in mind that the DTMF signals generated by telephone push button keypads are different from the Multi-Frequency (MF) tones generated by the telephone network to transmit information. You can use the DTMF Signals command to generate MF tones as well.

To generate DTMF signals:

1 Place the cursor where you want to insert the signals. Or, if you want to replace part of the existing waveform, select the desired range of audio data.

2 Choose Generate > DTMF Signals.

3 Set the following options as desired, and click OK:

Dial String Specifies the phone number for which you want to generate tones. You can also enter other characters, such as the asterisk (*) and pound (#) symbols, as well as the letters "A," "B," "C," and "D." Entering the pause character (see "Pause Character" in this list) inserts a pause of a defined length.

Tone Time Specifies the milliseconds for which the tones will last. The standard time for DTMF tones is 100 milliseconds.

Break Time Specifies the number of milliseconds of silence between successive tones.

Pause Time Specifies the length that is assigned to the pause character when it is used in the Dial String text box.

Pause Character Specifies which character Adobe Audition interprets as a pause.

DTMF Signals Generates DTMF signals by using combinations of the frequencies 697 Hz, 770 Hz, 852 Hz, 941 Hz and 1209 Hz, 1336 Hz, 1477 Hz, and 1633Hz.

MF Signals (CCITT R1) Generates MF signals (tones that are internal to telephone networks) using paired combinations of the frequencies 700 Hz, 900 Hz, 1100 Hz, 1300 Hz, 1500 Hz, and 1700 Hz.

Custom Specifies the combinations of frequencies to be used in generating signals. Select this option, and then enter values in the Hz text boxes of the keypad.

Amplitude Determines the volume level (as a percentage) of the tones generated, where 100% means maximum volume without clipping.

Twist Specifies how much louder the higher frequency tone is from the lower frequency tone. Enter a value (in decibels) in the Twist text box to increase the volume of the higher frequency tone accordingly.

Reset To DTMF Clears any custom frequency entries and replaces them with the standard DTMF frequency combinations.

Generating noise

The Noise command lets you generate random noise in a variety of colors. (Traditionally, *color* is used to describe the spectral composition of noise. Each color has its own characteristics.) Generating noise is useful for creating soothing sounds like waterfalls (perfect for use with the Binaural Auto-Panner function of Adobe Audition) and for generating signals that can be used to check out the frequency response of a speaker, microphone, or other audio system component.

To generate noise:

1 Place the cursor where you want to insert the noise. Or, if you want to replace part of the existing waveform, select the desired range of audio data.

2 Choose Generate > Noise.

3 Set the following options as desired, and click OK:

Color Specifies a color for the noise:

- Brown noise has a spectral frequency of $1/f^2$, which means, in layman's terms, that the noise has much more low-frequency content. Its sounds are thunder- and waterfall-like. Brown noise is so called because, when viewed, the wave follows a Brownian motion curve. That is, the next sample in the waveform is equal to the previous sample, plus a small random amount. When graphed, this waveform looks like a mountain range.
- Pink noise has a spectral frequency of $1/f$ and is found mostly in nature. It is the most natural sounding of the noises. By equalizing the sounds, you can generate rainfall, waterfalls, wind, rushing river, and other natural sounds. Pink noise is exactly between brown and white noise (hence, some people used to call it tan noise). It is neither random nor predictable; it is fractal-like when viewed. When zoomed in, the pattern looks identical to when zoomed out, except at a lower amplitude.
- White noise has a spectral frequency of 1, meaning that equal proportions of all frequencies are present. Because the human ear is more susceptible to high frequencies, white noise sounds very hissy. Adobe Audition generates white noise by choosing random values for each sample.

Style Specifies a style for the noise:

- Spatial Stereo generates noise by using three unique noise sources and spatially encoding them to seem as if one comes from the left, one from the center, and one from the right. When you listen to the result with stereo headphones, your mind perceives sound coming from all around. To specify the distance from center of the left and right noise sources, enter a delay value in microseconds. About 900 to 1000 microseconds correspond to the maximum delay perceivable. A delay of zero is identical to monaural noise, where left and right channels are the same.
- Independent Channels generates noise by using two unique noise sources, one for each channel. The left channel's noise is completely independent of the right channel's noise.
- Mono generates noise by using a single noise source, with the left and right channels set equally to that source.
- Inverse generates noise by using a single noise source (similar to the Mono option). However, the left channel's noise is exactly inverse of the right channel's noise. When you listen to the result with stereo headphones, your mind perceives sound coming from within your head instead of from somewhere externally.

Intensity Specifies the intensity of the noise on a scale of 2 to 40. At higher intensities, the noise becomes more erratic and sounds harsher and louder.

Duration Determines the number of seconds of noise that Adobe Audition generates.



For very long periods of noise, it's faster to generate a shorter period (say, about 10 to 20 seconds) and delete excess noise at the beginning and end so that the waves start and end at the midpoint. Then, copy and loop (choose Edit > Mix Paste) as many times as needed.

Generating tones

The Tones command lets you create a simple waveform and gives you control over numerous amplitude- and frequency-related settings. Generating tones is a great place to start when you create new sound effects.

To generate tones:

- 1 Place the cursor where you want to insert the tones. Or, if you want to replace part of the existing waveform, select the desired range of audio data.
- 2 Choose Generate > Tones.
- 3 Do one of the following:
 - To create a constant tone, select Lock To These Settings Only. Then, set options as desired, and click OK.
 - To create a tone that changes dynamically over time, deselect Lock To These Settings Only. Use the Initial Settings tab to set options for the initial tone, and use the Final Settings tab to set options for the final tone. After you set options, click OK. The tone generated will gradually go from the initial state to the final state.

For more information, see [“Generate Tones options” on page 92](#).

Generate Tones options

The Generate Tones dialog box provides the following options:

Base Frequency Specifies the main frequency to be used for generating tones.

Modulate By Modulates the base frequency in pitch over a user-defined range. For example, a 100 Hz setting modulates the original frequency by ± 100 Hz. At a setting of 100 Hz, a 1000 Hz tone would modulate between 900 Hz and 1100 Hz.

Modulation Frequency Specifies the rate (times per second) at which the frequency modulates. Entering a value of 10, for instance, generates tones that seem to warble in amplitude at the rate of 10 times per second. The tones actually warble in pitch (frequency) as they should, but because of the variance in the perceived energy levels of different frequencies to the human ear, they seem to warble in amplitude.

Frequency Components Adds up to five overtones to the fundamental frequency (Base Frequency).

Enter a multiplier for each overtone below the Frequency Components sliders. (The actual frequency will be this many times the fundamental.) Then, use the sliders to mix each of the individual components (0 to 100%) in proportion to one another. The overall gain (signal level) may be adjusted with the dB Volume sliders.

If Lock To These Settings Only isn't selected, all of the values can change over the duration of the audio file, so that they morph from the initial to final settings.

dB Volume Specifies the overall gain for each of the right and left channels from -80 dB to 0 dB. You can control both channels independently when generating stereo tones.

Start Phase Specifies the starting location in the cycle that will be produced. If Start Phase is set to 0 degrees, waves will start at the baseline. If Start Phase is set to 90 degrees, the wave will start at full amplitude (generating a noticeable click as well). If you work in great detail with tones and need to have the phase “just so,” this option gives you that control.

Phase Difference Allows the left channel to be out of phase with the right channel. A value of 0 causes the channels to be completely in phase and a value of 180 causes them to be completely out of phase.

Change Rate Dynamically changes the relative phase between the two channels of a stereo audio file over time at a given rate. For example, if you enter 1 Hz, the phase difference will cycle through 360 degrees each second.

DC Offset Adds a constant DC (Direct Current) amplitude to the tone, centering the waveform by shifting it up or down by the specified percentage. For example, you can apply DC Offset to correct an incoming signal that has suffered electrical pollution from a strong adjacent current.

Flavor Specifies the type of waveform to use. Each flavor has a particular sound unique unto itself. Sine waveforms are fundamental, with no harmonics (pure tone). Triangle waveforms have odd harmonics with amplitude of 1 to itself (squared). Square waveforms have odd harmonics with amplitude of 1 to itself. Sawtooth waveforms have all harmonics with amplitude of 1 to itself.

Duration Specifies the length of the generated tone (in seconds). Use decimals for partial seconds. For example, enter .25 to generate tones for exactly one-fourth of a second.

Modulate When a range of audio data is selected, causes the audio to be *ring modulated*, or multiplied, by the current tone settings. This option is great for adding special effects.

Demodulate When a range of audio data is selected, causes the audio to be demodulated. Use Demodulate on a previously modulated source to produce interesting effects.

Overlap (mix) When a range of audio data is selected, mixes the generated tones on top of the selected audio.

Note: For information on using presets and previewing results, search for [“Working with effects” on page 33](#).

Converting the sample type

A file's sample type determines its sample rate and bit depth, as well as the channel format (whether the waveform is mono or stereo). You can convert the sample type to change any of these attributes.

When you convert the sample type of a file, Adobe Audition directly processes the samples within the file, or resamples the data, so that the audio retains the same pitch and duration as the original file.

About sample rates

During the sampling process, an incoming analog signal is sampled at discrete time intervals. Each interval of analog signal is momentarily observed, and thus, each represents a specific, measurable voltage level. A mathematical conversion generates a digital series of numbers that represent the signal level at that particular point in time. The generated data can be digitally stored or processed.

The *sample rate* is the number of samples (or snapshots) that are taken of an audio signal per second. For example, a sample rate of 44,100 Hz means that 44,100 samples are taken per second. Since sampling is tied directly to the component of time, a system's sample rate determines a system's overall bandwidth—in other words, how many frequencies can be encoded within the audio signal. Higher sample rates generally yield a better quality waveform.

The most common sample rates for digital audio editing are as follows:

- 11,025 Hz Poor AM Radio Quality/Speech (low-end multimedia)
- 22,050 Hz Near FM Radio Quality (high-end multimedia)
- 32,000 Hz Better than FM Radio Quality (standard broadcast rate)
- 44,100 Hz CD Quality
- 48,000 Hz DAT Quality
- 96,000 Hz DVD Quality

Previewing a different sample rate

The Adjust Sample Rate command lets you preview how an audio file will sound at a different sample rate. This command doesn't convert the sample rate of the audio file—use the Convert Sample Type command to do that. (See [“Changing the sample rate” on page 94.](#))

To adjust the sample rate:


- 1 Choose Edit > Adjust Sample Rate.
- 2 Enter a sample rate in the text box, or choose a common sample rate from the list.
- 3 Click OK.

Note: Although you can create and edit any sample rate in Adobe Audition, your sound card may not be capable of playing it properly. To check the capabilities of your sound card, choose Options > Device Properties. (See [“Setting properties for audio output devices” on page 38.](#))

Changing the sample rate

The sample rate of a file determines the overall bandwidth of the waveform (that is, how many frequencies can be encoded within the audio signal). When changing the sample rate, keep in mind that most sound cards support only certain sample rates.

To change the sample rate of a file:

- 1 Choose Edit > Convert Sample Type. Alternatively, click the Convert Sample Type button  in the toolbar.
- 2 Select a rate from the Sample Rate list, or enter a custom rate in the text box.
- 3 Drag the Low/High Quality slider to adjust the quality of the sampling conversion.

Higher values retain more high frequencies (they prevent aliasing of higher frequencies to lower ones), but the conversion takes longer. Lower values requires less processing time but result in certain high frequencies being “rolled off,” leading to muffled-sounding audio. Usually, values between 100 and 400 are fine for most conversion needs.




Use higher values whenever you downsample a high rate to a low rate. When upsampling, results from lower values sound almost identical to those from higher values.

- 4 Select Pre/Post Filter to prevent false frequencies from being generated at the low end of the audio spectrum. Select this option for the best results.
- 5 Click OK.

Converting between stereo and mono

The Convert Sample Type command is the quickest way to convert a mono waveform into a stereo waveform, and vice versa. (You can also copy the waveform at its current volume directly into one channel or the other.) If you want to place separate waveforms on each channel of a stereo file and mix them at different volume levels, you can use the Mix Paste command instead.

To convert a waveform from mono to stereo, or vice versa:

- 1 Choose Edit > Convert Sample Type. Alternatively, click the Convert Sample Type button  in the toolbar.
- 2 Select Mono or Stereo.
- 3 Enter percentages for Left Mix and Right Mix:
 - When you convert a waveform from mono to stereo, the Left Mix and Right Mix options let you specify the relative amplitude with which the original mono signal is placed into each side of the new stereo signal. For example, you can place the mono source on the left channel only, the right channel only, or any balance point in between.
 - When you convert from stereo to mono, the Left Mix and Right Mix options let you control the amount of signal from the respective channel that will be mixed into the final mono waveform. The most common mixing method is to use 50% of both channels.
- 4 Click OK.



To remove all or most of the lead vocals from many stereo music recordings, you can convert a stereo waveform to mono with a Left Mix of 100% and a Right Mix of -100%. Most vocal tracks are positioned in the middle of the stereo field in-phase, so converting the signal so that it's out of phase often greatly reduces or eliminates the vocal track's level.

To create a stereo waveform with different waveforms in each channel:

- 1 Copy the mono waveform you want to place in the left channel.
- 2 Create a new file, and choose Edit > Mix Paste.
- 3 Select Overlap, and deselect Lock L/R. Set the left volume to 100%, set the right volume to 0%, and click OK.
- 4 Copy the mono waveform you want to place in the right channel.
- 5 Switch back to the new file you just created, and choose Edit > Mix Paste.
- 6 This time, set the left volume to 0% and the right volume to 100%. Click OK.

Changing the bit depth


The bit depth of a file determines the dynamic range of the audio. For example, 8-bit resolution provides 256 possible unique volumes, while 16-bit resolution provides 65,536 possible unique volumes. Adobe Audition supports up to 32-bit resolution.

You can raise the bit depth of a file to gain a greater dynamic range, or you can lower the bit depth to reduce the file size. When converting to a lower bit depth, Adobe Audition provides dithering options to help reduce noise and distortion. Although dithering introduces a small amount of white noise, the result is far preferable to the increased distortion that you would otherwise hear at low signal levels. Dithering also lets you hear sounds that would otherwise be masked by the noise and distortion limits of 8-bit audio.



Work at the 32-bit level when processing audio, even if you plan to downsample to 16- or 8-bit for output. You'll achieve better results than at the 16- or 8-bit level. The only time you may want to work at the 16- or 8-bit level is when processing a very large file on a slow computer.

To change the bit depth of a file:

- 1 Choose Edit > Convert Sample Type. Alternatively, click the Convert Sample Type button  in the toolbar.
- 2 Select a bit depth from the Resolution list, or enter a custom bit depth in the text box.
- 3 When you select a lower bit depth, options in the Dither section are enabled. Set the following options as desired, and click OK:

Enable Dithering Enables or disables dithering. If dithering is enabled, Adobe Audition truncates the audio, meaning that unused bits are simply chopped off and discarded. The result gives a crackly effect that fades in and out on very quiet audio passages.

Dither Depth (Bits) Sets the bit amount of dithering to be applied. In general, values of 0.2 to 0.7 give the best results without adding too much noise. Note, however, that as this value is lowered, other unwanted harmonic distortion noise appears. (Lower values are usually okay if you also apply Noise Shaping.)

p.d.f. (probability distribution function) Controls how the dithered noise is distributed away from the original audio sample value.

Usually, Triangular p.d.f. is a wise choice because it gives the best tradeoff among SNR (Signal-to-Noise ratio), distortion, and noise modulation. Triangular p.d.f. chooses random numbers that are generally closer to 0 than to the edges -1 or +1 (that is, the chance of 0 being chosen is twice as great as the chance of 0.5 or -0.5).

p.d.f.	SNR loss	Modulation
Rectangular	3 dB	Yes
Triangular	4.8 dB	No
Gaussian	6 dB	Negligible
Shaped Triangular	4.8 dB	No
Shaped Gaussian	6 dB	Negligible

Noise Shaping Determines the placement when you move noise to different frequencies. The same amount of overall noise is present, but you can place less noise at one frequency at the expense of placing more noise at another. You may also specify that no noise shaping is used.

Different curves result in different types of background noise. The type of curve to use depends on the source audio, final sample rate, and bit depth. By introducing noise shaping, you may be able to get away with a lower dither depths to reduce the overall background noise level, without introducing a lot of unwanted harmonic noise.


Curve	Sample Rate
Noise Shaping A	44.1 kHz or 48 kHz
Noise Shaping B	44.1 kHz or 48 kHz
Noise Shaping C1	44.1 kHz or 48 kHz
Noise Shaping C2	44.1 kHz or 48 kHz
Noise Shaping C3	44.1 kHz or 48 kHz
Noise Shaping D	44.1 kHz or 48 kHz
Noise Shaping E	44.1 kHz or 48 kHz
Noise Shaping E2	44.1 kHz or 48 kHz
Noise Shaping (44.1KHZ)	44.1 kHz
Noise Shaping (48KHZ)	48 kHz
Noise Shaping (96KHZ)	96 kHz

Note: In general, there are no really good noise shaping curves for audio at 32 kHz or lower. With audio at those sampling frequencies, try the different curves to see if they help, and just choose the one that sounds the best.

Converting multiple files to the same sample rate

If you need to make the same conversion on multiple files, you can save time by creating a sample rate conversion preset.

To create a sample rate conversion preset:

- 1 Choose Edit > Convert Sample Type. Alternatively, click the Convert Sample Type button  in the toolbar.
- 2 Adjust the settings as desired.
- 3 Click Save As, type a name for the preset, and click OK.

To apply a sample rate conversion preset:

Choose a preset from the list. The sample type settings change to the settings defined in the preset.

To delete a sample rate conversion preset:


Choose the preset from the list, and click Delete.

Adding file properties

The Wave Properties command opens a tabbed window that lets you add and get information about the active waveform.

Note: To preserve file properties when you save a file, make sure that you select *Save Extra Non-Audio Information*.

To add file information:

- 1 Choose View > Wave Properties. Alternatively, click the Add Information button  in the toolbar.
- 2 Click the tabs at the top of the dialog box to navigate between different sets of properties.
- 3 Set the properties as desired, and click OK.

Text Fields options

In Adobe Audition, you can embed text-based data in Windows .wav files that use the RIFF LIST INFO and DISP type 1 formats, and you can embed standard ID3 tag information in .mp3 files. Provided that other audio editors support this information, it remains with an audio file throughout its lifetime.

The options in the Text Fields tab depend on the setting you choose for Text Field Names: Standard RIFF, Radio Industry, or MP3 (ID3 Tag). Radio Industry format accommodates information for commercials and other types of audio files used by radio broadcasters. MP3 (ID3 Tag) format allows you to tag the opened .mp3 file with ID3v1.1-compatible data:

EBU Extensions options

The EBU Extensions tab of the Wave Properties dialog box provides options supported by the European Broadcasting Union.

Description Describes the audio file in up to 256 characters. You must enter a description before you can edit the other options.

Originator Specifies the name of the audio file's producer in up to 32 characters.

Originator Reference Specifies reference information about the producer in up to 32 characters.

Origination Date (yyyy-mm-dd) Specifies the date that the subject matter was produced. The date should be in the year-month-date format (yyyy-mm-dd). For example, specify June 8, 2004 as "2004-06-08."

Origination Time (hh:mm:ss) Specifies the time the audio file was produced. The format is hour:minutes:seconds, with the hour represented in Universal Military Time (for example, specify 10 p.m. as 22).

Time Reference (since midnight) Specifies the timecode of the audio file, calculated since 12:00 a.m. (midnight). Select from the hh:mm:ss.ddd or Samples options.

Coding History Provides a text box for you to describe all coding processes applied to the waveform. Adobe Audition doesn't automatically add any information in this text box; you must enter it manually.

Sampler options

The Sampler tab of the Wave Properties dialog box provides options relating to other devices, systems, or programs (such as synthesizer uploading and downloading software) that can be directly imbedded within .wav files. Provided that other audio editors support this information, it remains with a file throughout its lifetime.

Target Manufacturer ID Displays a proper value if a sampler has written a .wav file with this sampler information chunk present.

Target Product Code Displays a proper value if a sampler has written a .wav file with this sampler information chunk present.

Sample Period Specifies the sample rate of the file (or within 1 Hz of it). You can change the value in this text box if you wish the sampler to interpret the data at a different rate than it actually is.

Note Specifies a base (or root) note on a sampler that the current audio file is to be assigned to. The audio file's original pitch will be preserved whenever this key is played on a sampler.

Fine Tune Specifies the actual tone as a number of cents above the Note. You can enter values as precise as 1/100th of a cent.

Find Using Analysis Analyzes the audio file to determine automatically the Note and Fine Tune values. If a sampler loop is selected in the Sampler Loops list, the frequency at the center of that loop is entered into the Note and Fine Tune text boxes. If no loops are selected, the center of the entire waveform is used to gain the current note.

Note: *The Note and Fine Tune values can be off by a few hundredths of a cent, so you might need to adjust them manually after the note is found. For example, you might need to adjust G#4 at 99.99 cents to A4 at 0 cents.*

SMPTE Format Specifies the SMPTE frame rate format for the currently opened sample.

SMPTE Offset Specifies the SMPTE trigger offset point for the currently opened sample. For example, an audio file for a film soundtrack that needs to be triggered at 45 minutes, 15 seconds, and 29 frames might have a frame rate setting of 30 frames-per-second with an offset of 00:45:14:29.

Sampler Loops Lists sample loops. You can add new loops by first selecting a range of a waveform, and then clicking New in this tab. If no range is selected, click New to add the entire waveform as a new loop. You can also enter the actual starting point, ending point, and length in the appropriate text boxes.

Samplers can usually play loops forward, backward, or back and forth and back again. Each loop can be looped a different number of times or infinitely (as with a sustain loop, and the infinite loop would decay once the synth key is released). This information, however, is saved only in .wav files.

Misc options

The Misc tab of the Wave Properties dialog box lets you assign a .bmp or .dib image to an audio file. This image appears when you view the audio file's properties in Windows. For best results, choose a 32x32 pixel image.

If Use Default Wave Color is selected, then the color is the same as the one used for your current color scheme (the color of an unselected waveform in the Edit View wave display). If you want to select a different color, deselect Use Default Wave Color to open the Waveform Foreground Color dialog box, where you can select another color.

Cart options

Use the Cart tab of the Wave Properties dialog box to enter the Cart Chunk information for the file, if needed. The Cart Chunk is used by several popular radio automation packages. For details on how to use Cart Chunk data, see the documentation of your automation system.

File Info options

The File Info tab of the Wave Properties dialog box displays noneditable file information about the active audio file. Values include Filename, Folder, File Type, Uncompressed Size, File Format, Size On Disk, Date And Time Last Written, and Length.

Enhancing and Restoring Audio

About enhancing and restoring audio

If you need to add brilliance and impact to a new recording, or clean up the sound of an old one, you can use several types of audio enhancement and restoration effects:

- Noise reduction effects that let you remove unwanted hiss, hum, clicks, or pops. (See [“Removing noise” on page 108.](#))
- Filter effects that let you change overall tonal balance, from rumbling bass tones to sparkling highs. (See [“Filtering audio” on page 120.](#))
- Amplitude effects that let you precisely control audio volume for increased radio impact, detailed fade outs, and more. (See [“Optimizing amplitude” on page 130.](#))

All of these effects are available in Edit View, but some don't exist in Multitrack View. Because the two views are linked, however, you can easily overcome this limitation. If a multitrack clip requires noise reduction, for example, simply double-click the clip to process it in Edit View.

About the mastering process

Mastering describes the complete process of restoring and enhancing audio files for a particular medium, such as radio, video, CD, or the Web. In Adobe Audition, you can master either individual audio files in Edit View or groups of files in a batch process. (Batch processing is particularly useful if you plan to burn a group of files to CD. See [“About scripting and batch processing” on page 249.](#))

The mastering process consists of several stages, which are usually performed in the following order:

- 1. Analysis** To determine the overall frequency and dynamic range of the existing file. (See [“Analyzing frequency, phase, and dynamic range” on page 102.](#))
- 2. Noise reduction** To remove unwanted hiss, hum, clicks, or pops. (See [“Removing noise” on page 108.](#))
- 3. Equalization** To achieve the desired tonal balance. (See [“Filtering audio” on page 120.](#))
- 4. Compression** To maximize perceived volume. (See [“Optimizing amplitude” on page 130.](#))
- 5. Normalization** To ensure that the loudest sounds reach the highest possible level that digital systems allow—0 dBFS. (See [“Using the Normalize effect \(Edit View only\)” on page 137.](#))



You can reverse the order of the equalization and compression stages, but be aware that the volume of some tonal ranges may be over- or under-emphasized.



Before mastering audio, consider the requirements of the destination medium. If the destination is the Web, for example, the file will likely be played over speakers that poorly reproduce bass sounds. To compensate, you can boost bass frequencies during the equalization stage of the mastering process.

Analyzing frequency, phase, and dynamic range

In Edit View, you can analyze the frequency, phase, and dynamic range of an audio file. These analysis options are particularly helpful when used in conjunction with the many enhancement and restoration effects in Adobe Audition. For example, you can use the Frequency Analysis window to identify problematic frequency bands, which you can then correct with a filter effect. Similarly, you can use the Waveform Statistics dialog box to determine dynamic range and then compress that range with an amplitude effect.



To analyze a multitrack clip, double-click it to access Edit View.

Analyzing frequency range

In Edit View, you can use the Frequency Analysis window to analyze frequency range either statically for a selected area or dynamically during playback. In this window, the horizontal axis represents frequency (measured in Hz), while the vertical axis represents amplitude (measured in decibels).



To zoom in on a particular area of the Frequency Analysis graph, use the horizontal and vertical rulers. See [“Zooming graphs for frequency and phase analysis” on page 106](#).



Frequency Analysis window displaying Advanced options

A. Musical note **B.** Vertical ruler **C.** Horizontal ruler **D.** Left status area

E. Display options menu **F.** Right status area **G.** FFT type menu

To analyze frequency range:

- 1 In Edit View, select or play a range of the waveform.
- 2 Choose Window > Frequency Analysis, and set options as desired:

Linear View Sets the graph display to a linear horizontal frequency scale when selected or a logarithmic scale when deselected.

Hold buttons Take up to four frequency snapshots as a waveform is playing. The frequency outline (which is rendered in the same color as the button clicked) is frozen on the graph and overlaid on other frequency outlines. Up to four frozen frequency outlines may be shown at once. To clear a frozen frequency outline, click its corresponding Hold button again.

Status areas Display frequency and amplitude information directly underneath the graph. The left status area displays the highest frequency of the entire waveform and the maximum amplitude for each channel. The right status area displays the overall frequency (and equivalent musical note) at the center point of the selected range. The numbers beside musical notes indicate keyboard position and variance from standard tuning. For example, A2 +7 equals the second-lowest A on a keyboard tuned 7% higher than normal.



By default, the musical note of the left channel also appears at the top of the window. To hide that note, dock the window, right-click the window handle, and deselect Show Big Notes. For more information, see [“Using windows” on page 22](#).

Display style menu Select from the following graph display options:

- Lines displays amplitude at each frequency with simple lines. The left channel is blue; the right is red.
- Area (Left On Top) also displays lines for amplitude, but this option fills the area beneath the lines in a solid color, smooths out amplitude differences in the same area, and places the left channel in front.
- Area (Right On Top) functions identically to the option above, but places the right channel in front.
- Bars (Left On Top) shows the limitations on analysis resolution by splitting the display into rectangular segments, and places the left channel in front. The higher the FFT size, the greater the analysis resolution, and the narrower the bar.
- Bars (Right On Top) functions identically to the option above but places the right channel in front.

Scan Click this button to scan the highlighted selection and show all frequencies present in that selection.



By default, Adobe Audition analyzes only the center point of a selected range. To analyze the overall frequency of a selected range, click Scan.

For more information, see [“Advanced frequency analysis options” on page 104](#).

Advanced frequency analysis options

In the Frequency Analysis window, click Advanced to set the following options:

FFT Size menu Specifies the Fast Fourier Transform size. Higher FFT sizes give you more accurate results in terms of frequency (such as the overall frequency estimate), but they also create longer processing times.

You can generate a step-by-step animation by clicking the main waveform window and then holding down the Right Arrow key. As the cursor scrolls across the waveform, Adobe Audition displays the corresponding spectral information in the Frequency Analysis window.

Note: When you set the FFT size to 8192 or lower, the Frequency Analysis window updates in real time while you play a file. (Keep in mind that how well it ultimately updates in real time is based on the computer's speed.)

FFT type menu Lets you choose from eight types of FFT windows. Each displays a slightly different kind of frequency graph.

The Triangular window gives a more precise frequency estimate, but it's also the noisiest, meaning that other frequencies will be shown as present, even though they may be much lower in volume.

At the other extreme, the Blackmann-Harris window has a broader frequency band, which isn't as precise, but the sidelobes are very low, making it easier to pick out the major frequency components.

Reference Determines the amplitude at which full scale, 0 dBFS audio data is displayed. For example, a value of zero displays 0 dBFS audio at 0 dB. A value of 30 displays 0 dBFS audio at -30 dB. This value simply moves the display up or down; it does not change the amplitude of audio data.

Copy To Clipboard Copies a text-based frequency report of the current waveform to the Windows clipboard.

Analyzing phase


In Edit View, you can use the Phase Analysis window to analyze phase either statically for a selected range or dynamically during playback. You should analyze phase only for stereo waveforms, as phase differences don't exist in mono waveforms. Phase analysis can reveal out-of-phase channels, which you can correct with the Invert command. (See [“Inverting and reversing audio” on page 89.](#))

The Phase Analysis window includes a Lissajou Plot graph. By default, this graph displays phase differences between the left and right channels as follows:

- A mono waveform appears as a diagonal line ascending from left to right.
- A right-channel-only waveform appears as a horizontal line.
- A left-channel-only waveform appears as a vertical line.
- A completely out-of-phase stereo waveform appears as a diagonal line descending from left to right.
- A typical stereo waveform appears as many wavy lines descending from right to left.
- A stereo waveform with wide separation appears as many wavy lines extending in all directions.



To zoom in on a particular area of the Phase Analysis graph, use the horizontal and vertical rulers. See [“Zooming graphs for frequency and phase analysis” on page 106](#).



Phase Analysis window with display menu revealed

To analyze phase:

- 1 In Edit View, select or play a range of the waveform.
- 2 Choose Analyze > Show Phase Analysis, and set options as desired:

Normalize Enlarges the phase analysis lines so that they reach the edge of the graph.

Display menu Select from the following options:

- Left/Right to display the defaults noted in the introduction above.
- Mid/Side to rotate the display to the left by 45 degrees. The horizontal ruler (x-axis) plots the side channel $[(\text{right} - \text{left})/2]$ while the vertical ruler (y-axis) plots the mid channel $[(\text{right} + \text{left})/2]$.
- Spin to display the waveform on a phase graph rather than an amplitude graph.

Samples Defines the number of samples displayed concurrently. Higher sample sizes give you more accurate results, but they require much more processing power to be effective. Choose the sample size that best suits your system.

For more information, see [“Advanced phase analysis options” on page 106](#).

Advanced phase analysis options

When the Phase Analysis window is docked, you can right-click the window handle to set the following advanced options:

Allowing Drawing Adds a Draw button to the window. Click the Draw button to draw in the phase graph with the Pencil tool.


Note: *Drawing in the phase graph is destructive, so changes are permanent if you save the file.*

Spin Displays amplitude rather than phase, redrawing the waveform display on an axis that rotates in the phase graph.

Edit Spin Rate Determines how fast the spin axis rotates.

Zooming graphs for frequency and phase analysis

In the Frequency Analysis and Phase Analysis windows, you can zoom graphs to analyze frequency and phase in more detail.



Zooming and navigating a Phase Analysis graph

To zoom in on a graph:

In the vertical or horizontal ruler, right-click and drag the magnifying glass icon.

To navigate a magnified graph:

In the vertical or horizontal ruler, left-click and drag the hand icon .

To zoom out on a magnified graph:

Right-click in the vertical or horizontal ruler, and choose one of the following from the pop-up menu:


- Zoom Out to return to the previous magnification. (This option is available only in the Frequency Analysis window.)
- Zoom Out Full to zoom out completely.

Viewing waveform statistics

In Edit View, you can use the Waveform Statistics dialog box to evaluate a variety of information about audio amplitude. This dialog box contains two tabs, General and Histogram, both of which share an RMS Settings section. The General tab displays numerical text boxes that indicate dynamic range, identify clipped samples, and note any DC offset. The Histogram tab displays a graph that shows the relative prevalence of each amplitude: The horizontal ruler measures amplitude in decibels, and the vertical ruler measures prevalence using the RMS formula.



Use the Histogram tab to identify prevalent amplitudes, and then compress, limit, or normalize them with an amplitude effect. (See [“Optimizing amplitude” on page 130.](#))



Waveform Statistics dialog box, Histogram tab

To view a waveform histogram:

- 1 In Edit View, select an audio range.
- 2 Choose Analyze > Statistics, and click the Histogram tab.
- 3 Select Left or Right to display either the left or right channel in the foreground.

To view numerical waveform statistics:

- 1 In Edit View, select an audio range.
- 2 Choose Analyze > Statistics, and click the General tab.

For more information, see [“Waveform Statistics options” on page 107.](#)

Waveform Statistics options

The Waveform Statistics dialog box provides the following options:

Minimum Sample Value Shows the sample with the lowest amplitude. Click the arrow button → to place the current-time indicator at that location and close the Waveform Statistics dialog box.

Maximum Sample Value Shows the sample with the highest amplitude. Click the arrow button → to place the current-time indicator at that location and close the Waveform Statistics dialog box.

Peak Amplitude Shows the sample with the highest amplitude in decibel form. Click the arrow button → to place the current-time indicator at that location and close the Waveform Statistics dialog box.

Possibly Clipped Samples Shows the number of samples that could exceed 0 dBFS. Click the arrow button → to place the current-time indicator at the first clipped sample and close the Waveform Statistics dialog box. (If necessary, choose Analyze > Statistics, and click this arrow button again to detect subsequent clipped samples.)

DC Offset Shows the direct current offset of the center of the waveform, measured in percent. Positive values are above the center line (zero volts), and negative values are below it.

Minimum RMS Power Shows the minimum RMS amplitude. Click the arrow button → to place the current-time indicator at that location and close the Waveform Statistics dialog box.

Maximum RMS Power Shows the maximum RMS amplitude. Click the arrow button → to place the current-time indicator at that location and close the Waveform Statistics dialog box.

Average RMS Power Shows the average amplitude. This value reflects perceived loudness.

Total RMS Power Represents the total power of the entire selection.

Actual Bit Depth Reports the waveform's bit depth (or "float" if the waveform uses the full 32-bit float range).

Copy Data to Clipboard Copies all statistics on the General tab.

RMS settings Provides the following options:

- 0dB = FS Sine Wave sets the dB level of the RMS settings to correspond to a full-scale sine wave (where peak amplitude is at 0 dB, using every sample value in the 16-bit range).
- 0dB = FS Square Wave sets the dB level of the RMS settings to correspond to a full-scale square wave, where peak amplitude is about 3.02 dB louder than a full-scale sine wave.
- Account For DC subtracts any DC offset to achieve the most accurate RMS values.
- Window Width specifies the number of milliseconds in each RMS window. A selected range contains a series of such windows, which Adobe Audition averages to calculate the Minimum RMS and Maximum RMS values. To achieve the most accurate RMS values, use wide windows for audio with a wide dynamic range, and narrow windows for audio with a narrow dynamic range.
- Recalculate RMS updates the RMS values after you specify new RMS settings.

Removing noise

In Edit View, you can use effects in the Noise Reduction menu to reduce background noise and broadband noise without reducing audio quality.

Using the Auto Click/Pop Eliminator effect (Edit View only)

If you need to quickly remove crackle and static from vinyl recordings, first try the Auto Click/Pop Eliminator effect. You can easily select and correct a large area of audio or a single click or pop. This effect provides the same processing quality as the Click/Pop Eliminator effect, but it offers simplified controls and a helpful preview.

To use the Auto Click/Pop Eliminator effect:

- 1 In Edit View, select an audio range.
- 2 In the Effects tab of the Organizer window, expand Noise Reduction, and double-click Auto Click/Pop Eliminator.
- 3 Set the desired options.

For more information, see [“Auto Click/Pop Eliminator options” on page 109](#).

Auto Click/Pop Eliminator options

The Auto Click/Pop Eliminator effect provides the following options:

Noise Threshold Determines sensitivity to noise. Lower settings detect more clicks and pops but may include audio you wish to retain. Settings range from 1 to 100; the default is 35.


Complexity Indicates the complexity of noise. Higher settings apply more processing but can degrade audio quality. Settings range from 1 to 100; the default is 1.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Click/Pop Eliminator effect (Edit View only)

The Click/Pop Eliminator effect detects and removes clicks and pops. Like the Auto Click/Pop Eliminator, this effect is ideal if you want to clean up the sound of vinyl recordings before transferring them to CD or another digital medium. The Click/Pop Eliminator, however, provides a much wider range of controls, letting you highly customize settings for unique audio content.

For this effect, the most important parameters are the Detect and Reject thresholds. (To enable the latter, you must select Second Level Verification.) For Detect thresholds, try settings ranging from 10 for a lot of correction to 50 for very little correction. For Reject thresholds, try settings ranging from 5 to 40. Run Size is the second most important parameter. A setting of about 25 is best for high-quality work. For the highest quality, apply the Click/Pop Eliminator in three successive passes (where each pass is faster than the previous one).



Click/Pop detection graph

A. Level of detected clicks and pops **B.** Level of rejected clicks and pops



To visually identify clicks, zoom in and use Spectral View with a resolution of 256 bands and a window width of 40%. (You can access these settings in the Display tab of the Settings dialog box.) Most clicks appear as bright vertical bars that extend from the top to the bottom of the waveform display.

To use the Click/Pop Eliminator effect:

- 1 In Edit View, select an audio range.
- 2 In the Effects tab of the Organizer window, expand Noise Reduction, and double-click Click/Pop Eliminator.
- 3 Set the desired options.

For more information, see [“Click/Pop Eliminator options” on page 110](#).

Click/Pop Eliminator options

The Auto Click/Pop Eliminator effect provides the following options:

Detection graph Shows the exact threshold levels to be used at each amplitude, with amplitude along the horizontal ruler (x-axis) and threshold level along the vertical ruler (y-axis). Adobe Audition uses values on the curve to the right (above –20 dB or so) when processing louder audio and values on the left when processing softer audio. Curves are color-coded to indicate detection and rejection.

Auto Find All Levels Scans the highlighted area for clicks based on the values for Sensitivity and Discrimination, and determines values for Threshold, Detect, and Reject. Five areas of audio are selected, starting at the quietest and moving to the loudest.

Sensitivity Determines the level of clicks to detect. Use a lower value, such as 10, to detect lots of subtle clicks, or a value of 20 to detect a few louder clicks. (Detected levels with Auto Find All Levels are always higher than with this option.)

Discrimination Determines how many clicks to fix. Enter high values to fix very few clicks and leave most of the original audio intact. Enter lower values, such as 20 or 40, if the audio contains a moderate number of clicks. Enter extremely low values, such as 2 or 4, to fix constant clicks.

Find Threshold Levels Only Automatically sets the Max (Maximum), Avg (Average), and Min (Minimum) Threshold levels.

Max, Avg, and Min Thresholds Determine the unique detection and rejection thresholds for the maximum, average, and minimum amplitudes of the audio. For example, if audio has a maximum RMS amplitude of –10 dB, you should set Max Threshold to –10 dB. If the minimum RMS amplitude is –55 dB, then set Min Threshold to –55.

Set the threshold levels before you adjust the corresponding Detect and Reject values. (Set the Max and Min Threshold levels first, because once they're in place, you shouldn't need to adjust them much.) Set the Avg Threshold level to about three quarters of the way between the Max and Min Threshold levels. For example, if Max Threshold is set to 30 and Min Threshold is set to 10, set Avg Threshold to 25.

After you audition a small piece of repaired audio, you can adjust the settings as needed. For example, if a quiet part still has a lot of clicks, lower the Min Threshold level a bit. If a loud piece still has clicks, lower the Avg or Max Threshold level. In general, less correction is required for louder audio, as the audio itself masks many clicks, so repairing them isn't necessary. Clicks are very noticeable in very quiet audio, so quiet audio tends to require lower detection and rejection thresholds.

Detect Determines sensitivity to clicks and pops. Possible values range from 1 to 150, but recommended values range from 6 to 60. Lower values detect more clicks.

Start with a threshold of 35 for high-amplitude audio (above -15 dB), 25 for average amplitudes, and 10 for low-amplitude audio (below -50 dB). These settings allow for the most clicks to be found, and usually all of the louder ones. If a constant crackle is in the background of the source audio, try lowering the Min Threshold level or increasing the dB level to which the threshold is assigned. The level can be as low as 6, but a lower setting can cause the filter to remove sound other than clicks.

If more clicks are detected, more repair occurs, increasing the possibility of distortion. With too much distortion of this type, audio begins to sound flat and lifeless. If this occurs, set the detection threshold rather low, and select Second Level Verification to reanalyze the detected clicks and disregard percussive transients that aren't clicks.



If you still hear clicks after filtering audio, lower the detection threshold; if audio becomes too distorted, either increase the threshold or select Second Level Verification.

Reject Determines how many potential clicks (found using the Detection Threshold) are rejected if Second Level Verification box is selected. Values range from 1 to 100; a setting of 30 is a good starting point. Lower settings allow for more clicks to be repaired. Higher settings can prevent clicks from being repaired, as they might not be actual clicks.

You want to reject as many detected clicks as possible but still remove all audible clicks. If a trumpet-like sound has clicks in it, and the clicks aren't removed, try lowering the value to reject fewer potential clicks. If a particular sound becomes distorted, then increase the setting to keep repairs at a minimum. (The fewer repairs that are needed to get good results, the better.)

Second Level Verification Rejects some of the potential clicks found by the click detection algorithm. In some types of audio, such as trumpets, saxophones, female vocals, and snare drum hits, normal peaks are sometimes detected as clicks. If these peaks are corrected, the resulting audio will sound muffled. Second Level Verification rejects these audio peaks and corrects only true clicks



This option reduces performance, so you should use it only for sections that are very troublesome.

Pulse Train Verification Prevents normal waveform peaks from being detected as clicks. It may also reduce detection of valid clicks, requiring more aggressive threshold settings. Select this option only if you've already tried to clean up the audio but stubborn clicks remain.

Link Channels Analyzes audio from both channels simultaneously. If a click is found in one channel, a click will most likely be detected in the other.

Smooth Light Crackle Smooths out one-sample errors when detected, often removing more background crackle. If the resulting audio sounds thinner, flatter, or more tinny, deselect this option.

Detect Big Pops Removes large unwanted events (such as those more than a few hundred samples wide) that might not be detected as clicks. Values can range from 30 to 200.

Note that a sharp sound like a loud snare drum hit can have the same characteristic as a very large pop, so select this option only if you know the audio has very large pops (like a vinyl record with a very big scratch in it). If this option causes drum hits to sound softer, slightly increase the threshold to fix only loud, obvious pops.

If loud, obvious pops aren't fixed, select Detect Big Pops, and use settings from about 30 (to find quiet pops) to 70 (to find loud pops).

Multiple Passes Performs up to 32 passes automatically to catch clicks that might be too close together to be repaired effectively. Fewer passes occur if no more clicks are found and all detected clicks are repaired. In general, about half as many clicks are repaired on each successive pass. A higher detection threshold might lead to fewer repairs and increase the quality while still removing all clicks.

FFT Size Determines the FFT size used to repair clicks, pops, and crackle. In general, select Auto to let Adobe Audition determine the FFT size. For some types of audio, however, you might want to enter a specific FFT size (from 8 to 512). A good starting value is 32, but if clicks are still quite audible, increase the value to 48, and then 64, and so on. The higher the value, the slower the correction will be, but the better the potential results. If the value is too high, rumbly, low frequency distortion can occur.



If you repair clicks one at a time by clicking Fill Single Click Now, a high FFT size (128 to 256) works well.

Pop Oversamples Includes surrounding samples in detected clicks. When a potential click is found, its beginning and end points are marked as closely as possible. The Pop Oversamples value (which can range from 0 to 300) expands that range, so more samples to the left and right of the click are considered part of the click.

If corrected clicks become quieter but are still evident, increase the Pop Oversamples value. Start with a value of 8, and increase it slowly to as much as 30 or 40. Audio that doesn't contain a click shouldn't change very much if it's corrected, so this buffer area should remain mostly untouched by the replacement algorithm.

Increasing the Pop Oversamples value also forces larger FFT sizes to be used if Auto is selected. A larger setting may remove clicks more cleanly, but if it's too high, audio will start to distort where the clicks are removed.

Run Size Specifies the number of samples between separate clicks. Possible values range from 0 to 1000. To independently correct extremely close clicks, enter a low value; clicks that occur within the Run Size range are corrected together.

A good starting point is around 25 (or half the FFT size if Auto next to FFT Size isn't selected). If the Run Size value is too large (over 100 or so), then the corrections may become more noticeable, as very large blocks of data are repaired at once. If you set the Run Size too small, then clicks that are very close together may not be repaired completely on the first pass.

Fill Single Click Now Corrects a single click in a selected audio range. If Auto is selected next to FFT Size, then an appropriate FFT size is used for the restoration based on the size of the area being restored. Otherwise, settings of 128 to 256 work very well for filling in single clicks. Once a single click is filled, press the F3 key to repeat the action. You can also create a quick key in the Favorites menu for filling in single clicks.

Note: If the Fill Single Click Now button is unavailable, the selected audio range is too long. Click Cancel, and select a shorter range in the waveform display.

Corrected and Rejected Indicates how many clicks were corrected, in addition to how many rejected clicks that would have been corrected if Second Level Verification wasn't enabled.



To hear all the clicks that were removed, use the Mix Paste command to combine a copy of the original file with the corrected audio. Choose Edit > Mix Paste, select the original file, enter 100% for Volume, and select Invert. For more information, see ["Mixing audio data when pasting" on page 80](#).

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see ["Working with effects" on page 33](#).

Using the Clip Restoration effect (Edit View only)

The Clip Restoration effect repairs clipped waveforms by filling in clipped sections with new audio data. Clipping occurs when the amplitude of a signal exceeds the maximum level for the current bit resolution (for example, levels above 256 in 8-bit audio). Commonly, clipping results from recording levels that are too high. You can monitor clipping during recording or playback by watching the Level Meters; when clipping occurs, the boxes on the far right of the meters turn red.

Visually, clipped audio appears as broad flat areas at the top of a waveform. Sonically, clipped audio is a static-like distortion.

Note: If you need to adjust the DC offset of clipped audio, first use the Clip Restoration effect. If you instead adjust DC offset first, the Clip Restoration effect won't identify clipped areas that fall below 0 dBFS.

To restore clipped audio:

- 1 In Edit View, select an audio range.
- 2 In the Effects tab of the Organizer window, expand Noise Reduction, and double-click Clip Restoration.
- 3 Set the desired options.

For more information, see ["Clip Restoration options" on page 113](#).

Clip Restoration options

The Clip Restoration effect provides the following options:

Input Attenuation Specifies the amount of amplification that occurs before processing.

Overhead Specifies the percentage of variation in clipped regions. A value of 0% detects clipping only in perfectly horizontal lines at maximum amplitude. A value of 1% detects clipping beginning at 1% below maximum amplitude. (A value of 1% detects almost all clipping and leads to a more thorough repair.)

Minimum Run Size Specifies the length of the shortest run of clipped samples to repair. A value of 1 repairs all samples that seem to be clipped, while a value of 2 repairs a clipped sample only if it's followed or preceded by another clipped sample.

FFT Size Sets an FFT (Fast Fourier Transform) Size, measured in samples, if audio is severely clipped (for example, because of too much bass). In this case, you want to estimate the higher frequency signals in the clipped areas. Using the FFT Size option in other situations might help with some types of clipping. (Try a setting of 40 for normal clipped audio.) In general, however, leave FFT Size unselected. If FFT Size is unselected, Adobe Audition uses spline curve estimation.

Clipping Statistics Shows the minimum and maximum sample values found in the current selected range, as well as the percent of samples clipped based on that data.

Gather Statistics Now Updates the Clipping Statistics values for the current selection or file.




To retain amplitude when restoring clipped audio, work at 32-bit resolution for more precise editing. (See [“Changing the bit depth” on page 96](#).) Then apply the Clip Restoration effect with no attenuation, followed by the Hard Limiting effect with a Boost value of 0 and a Limit value of -0.2 dB.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Hiss Reduction effect (Edit View only)

The Hiss Reduction effect reduces hiss from sources such as audio cassettes, vinyl records, or microphones. This effect greatly lowers the amplitude of a frequency range if it falls below an amplitude threshold called the *noise floor*. Audio in frequency ranges that are louder than the threshold remain untouched. If audio has a consistent level of background hiss, that hiss can be removed completely.



Using the Hiss Reduction graph to adjust the noise floor



To reduce other types of noise that have a wide frequency range, try the Noise Reduction effect. (See [“Using the Noise Reduction effect \(Edit View only\)” on page 116](#).)

To reduce hiss:

- 1 In Edit View, select an audio range.
- 2 In the Effects tab of the Organizer window, expand Noise Reduction, and double-click Hiss Reduction.
- 3 Set the desired options.

For more information, see [“Hiss Reduction options” on page 115](#).

Hiss Reduction options

The Hiss Reduction effect provides the following options:

Graph Represents the estimated noise floor that exists for each frequency in the source audio, with frequency along the horizontal ruler (x-axis) and amplitude, or noise floor, along the vertical ruler (y-axis). This information helps to distinguish hiss from desirable audio data.

The actual value used to perform hiss reduction is a combination of the graph and the Noise Floor Adjust slider, which shifts the estimated noise floor reading up or down for fine tuning.

Viewing Displays the waveform's left or right channel.

Get Noise Floor Graphs an estimate of the noise floor. The estimate is used by the Hiss Reduction effect to more effectively remove just hiss while leaving regular audio untouched. Get Noise Floor is the most powerful feature of Hiss Reduction.

To create a graph that most accurately reflects the noise floor, click Get Noise Floor with a section of just hiss highlighted in a waveform. If you can't identify such a section, select an area that has the least amount of music or other desired audio, in addition to the least amount of high frequency information. (In spectral view, such an area won't have any activity in the top 75% of the display.)

After you identify the noise floor, you might need to lower the control points on the left (representing the lower frequencies) to make the graph as flat as possible. If music was present at any frequency, the control points around that frequency will be higher than they should be.

Drag Points Specifies the number of drag points, or control points, on the graph.

Reset (Hi, Med, Low) Resets the estimated noise floor. Click Hi to set the noise floor at – 50 dB (for very loud hiss), click Med to set the floor to –70 dB (for average hiss), or click Low to set the floor to –90 dB (for very little hiss).



For quick, general-purpose hiss reduction, a complete noise floor graph isn't always necessary. In many cases, you can simply reset the graph to an even level and manipulate the Noise Floor Adjust slider.

Noise Floor Adjust Fine tunes the noise floor until the appropriate amount of hiss reduction and quality level is achieved.

FFT Size Specifies a transform size. In general, sizes from 3000 to 6000 work best.

- Lower FFT sizes (2048 and below) result in better time response (less swooshing before cymbal hits, for example), but they can produce poorer frequency resolution, creating hollow or flanged sounds.
- Higher FFT sizes (12,000 and above) might cause swooshing, reverb, and drawn out background tones, but they produce very accurate frequency resolution.

Precision Factor Determines the accuracy of hiss reduction in the time domain and affects the decay rate of spectral components below the previous hiss level. (See Spectral Decay Rate.) Typical values range from 7 to 14.

- Larger values generally produce better results and slower processing speeds. Values over 20 don't ordinarily improve quality any further.
- Lower values might result in a few milliseconds of hiss before and after the louder parts of audio.

Transition Width Produces a slow transition in hiss reduction instead of an abrupt change from no reduction to the reduced hiss level. Values from 5 to 10 usually achieve good results.

- If the value is too low, other background artifacts might be heard.
- If the value is too high, some hiss may remain after processing.

Spectral Decay Rate When audio is encountered above the estimated noise floor, determines how much audio in the same frequency band is assumed to follow. With low values, less audio is assumed to follow, and the carving function will cut more closely in time to the frequencies being kept. Values of 40% to 75% work best.

- If the value is too low, background bubbly effects might be heard, and music might sound artificial.
- If the value is too high (above 90%), unnaturally long tails and reverbs might be heard.

Reduce Hiss By Sets the level of hiss reduction for audio that is below the estimated noise floor.

- With lower values, not as much noise is removed, and the original audio signal stays relatively undisturbed.
- With higher values (especially above 20 dB) dramatic hiss reduction can be achieved, but the remaining audio might become distorted.

Remove Hiss, Keep Only Hiss: Removes hiss or removes all audio except for hiss.


Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Noise Reduction effect (Edit View only)

The Noise Reduction effect dramatically reduces background and broadband noise with a minimal reduction in signal quality. This effect can remove a wide range of noise, including tape hiss, microphone background noise, 60-cycle hum, or any noise that is constant throughout a waveform.

The proper amount of noise reduction depends upon the type of background noise and the acceptable loss in quality for the remaining signal. In general, you can increase the signal-to-noise ratio by 5 to 20 dB and retain high audio quality.

To achieve the best results with the Noise Reduction effect, apply it to 16- or 32-bit audio with no DC offset. With 8-bit audio, this effect cannot reduce noise below – 45 dB, which is very audible. (To achieve a lower noise floor with 8-bit audio, upsample the file to 16 bits, apply the Noise Reduction effect, and downsample the file back to 8 bits.) With a DC offset, this effect may introduce clicks in quiet passages. (To remove a DC offset, select the Center Wave preset provided by the Amplify/Fade effect.)



Adjusting frequency-specific settings with the Noise Reduction graphs:
A. Noise floor **B.** Reduction graph **C.** Original audio **D.** Processed audio



To reduce noise added by a sound card during recording, start the recording with a second of silence. After recording is complete, use that silence as the Noise Reduction Profile, and then remove it from the complete recording. In some cases, this process can increase dynamic range by 10 dB.

To reduce noise:

1 In Edit View, select a range that contains only noise and is at least half a second long.



To select noise in a specific frequency range, use the Marquee Selection tool. (See [“Selecting audio frequencies in Spectral View” on page 77.](#))

2 In the Effects tab of the Organizer window, expand Noise Reduction, and double-click Capture Noise Reduction Profile.

3 In the waveform display, select the range from which you want to remove noise.

4 In the Effects tab of the Organizer window, double-click Noise Reduction.

5 Set the desired options.

For more information, see [“Noise Reduction options” on page 117.](#)

Noise Reduction options

The Noise Reduction effect provides the following options:

View Displays either the left or right channel noise profile. The amount of noise reduction is always the same for both channels. To perform separate levels of reduction on each channel, edit the channels individually.

Noise Profile graph Represents, in yellow, the amount of noise reduction that will occur at any particular frequency. Adjust the graph by moving the Noise Reduction Level slider.

Capture Profile Extracts a noise profile from a selected range, indicating only background noise. Adobe Audition gathers statistical information about the background noise so it can remove it from the remainder of the waveform.



If the selected range is too short, Capture Profile is disabled. Reduce the FFT Size or select a longer range of noise. If you can't find a longer range, copy and paste the currently selected range to create one. (You can later remove the pasted noise by using the Edit > Delete Selection command.)

Snapshots In Profile Determines how many snapshots of noise to include in the captured profile. A value of 4000 is optimal for producing accurate data.

Very small values greatly affect the quality of the various noise reduction levels. With more samples, a noise reduction level of 100 will likely cut out more noise, but also cut out more original signal. However, a low noise reduction level with more samples will also cut out more noise, but likely will not disrupt the intended signal.

Load From File Opens a previously saved noise profile. You can open any .fft file that Adobe Audition has saved.

Note: A noise profile can be used only on a sample of the same type. In other words, a 22 kHz, mono, 16-bit profile can't be used with a 44K, stereo, 8-bit sample. Bear in mind also that because noise profiles are so specific, a profile for one type of noise is likely to not produce good results when used for another type of noise—regardless of whether the sample types are compatible. Even if audio samples were recorded with the same microphone, the type of background noise could be different if the recording environment is different.

Save Saves the noise profile as an .fft file, which contains information about sample type, FFT (Fast Fourier Transform) size, and three sets of FFT coefficients: one for the lowest amount of noise found, one for the highest amount, and one for the power average.

Select Entire File Lets you apply a previously captured noise-reduction profile to the entire file.

Reduction graph Sets the amount of noise reduction at certain frequency ranges. For example, if you need noise reduction only in the higher frequencies, adjust the chart to give less noise reduction in the low frequencies, or alternatively, more reduction in the higher frequencies.

The graph depicts frequency along the x-axis (horizontal) and the amount of noise reduction along the y-axis (vertical). If the graph is flattened (click Flat), then the amount of noise reduction used is based on the noise profile exactly. The readout below the graph displays the frequency and adjustment percentage at the position of the cursor.

Log Scale Displays the Noise Profile graph in either linear or logarithmic fashion.

- Select Log Scale to divide the graph evenly into 10 octaves.
- Deselect Log Scale to divide the graph linearly, with each 1000 kHz (for example) taking up the same amount of horizontal width.

Live Update Enables the Noise Profile graph to be redrawn as you move control points on the Reduction graph.

Noise Reduction Level Adjusts the amount of noise reduction to be applied to the waveform or selection. Alternatively, enter the desired amount in the text box to the right of the slider.

Note: Depending on the original waveform and the type of noise removed, high noise reduction levels can sometimes cause the remaining audio to have a flanged or phaselike quality. For better result, undo the effect and try a lower setting.

Noise Reduction Settings Provides the following options:

- **FFT Size:** Determines how many individual frequency bands are analyzed. This option causes the most drastic changes in quality. The noise in each frequency band is treated separately, so the more bands you have, the finer frequency detail you get in removing noise. For example, if there's a 120 Hz hum, but not many frequency bands, frequencies from 80 Hz on up to 160 Hz may be affected. With more bands, less spacing between them occurs, so the actual noise is detected and removed more precisely. However, with too many bands, time slurring occurs, making the resulting sound reverberant or echo-like (with pre- and post-echoes). So the tradeoff is frequency resolution versus time resolution, with lower FFT sizes giving better time resolution and higher FFT sizes giving better frequency resolution. Good settings for FFT Size range from 4096 to 12000.
- **Remove Noise, Keep Only Noise:** Removes noise or removes all audio except for noise.
- **Reduce By:** Aids in reducing bubbly background effects. Values between 5 and 100 dB work well.
- **Precision Factor:** Affects distortions in amplitude. Values of 5 and up work best, and odd numbers are best for symmetric properties. With values of 3 or less, the FFT is performed in giant blocks and a drop or spike in volume can occur at the intervals between blocks. Values beyond 10 cause no noticeable change in quality, but they increase the processing time.
- **Smoothing Amount:** Takes into account the standard deviation, or variance, of the noise signal at each band. Bands that vary greatly when analyzed (such as white noise) will be smoothed differently than constant bands (like a 60 cycle hum). In general, increasing the smoothing amount (up to 2 or so) reduces burbly background artifacts at the expense of raising the overall background broadband noise level.
- **Transition Width:** Determines the range between what is noise and what remains. For example, a transition width of zero applies a sharp, noise gate-type curve to each frequency band. If the audio in the band is just above the threshold, it remains; if it's just below, it's truncated to silence. Conversely, you can specify a range over which the audio fades to silence based upon the input level. For example, if the transition width is 10 dB, and the cutoff point (scanned noise level for the particular band) is -60 dB, then audio at -60 dB stays the same, audio at -62 dB is reduced (to about -64 dB), and so on, and audio at -70 dB is removed entirely. Again, if the width is zero, then audio just below -60 dB is removed entirely, while audio just above it remains untouched. Negative widths go above the cutoff point, so in the preceding example, a width of -10 dB creates a range from -60 to -50 dB.

- **Spectral Decay Rate:** Specifies the percentage of frequencies processed when audio falls below the noise floor. Fine-tuning this percentage allows greater noise reduction with fewer artifacts. Values of 40% to 75% work best. Below those values, bubbly-sounding artifacts are often heard; above those values, excessive noise typically remains.



You can create unusual effects by using foreground audio as a noise profile rather than background noise. In a vocal recording, for example, you can use the vowel sound “oh” as the profile and then reduce or eliminate “oh” sounds throughout the recording.


Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Filtering audio

Filter effects change the frequency content of audio, letting you adjust tonal range to enhance audio or create special effects. (Be aware, however, that significantly boosting a frequency can cause clipping.)

Using the Dynamic EQ effect

The Dynamic EQ effect varies the amount of equalization over time. For example, during the first half of a waveform, you can boost high frequencies; during the second half, you can change the bandwidth of affected frequencies. The Dynamic EQ dialog box provides three areas of controls: Gain, Frequency, and Q (bandwidth).



Frequency graph of the Dynamic EQ effect in Edit View (Rhythmic Sweep preset)



Dynamic EQ is especially effective as a real-time effect in Multitrack View, where you can use clip envelopes to adjust the Gain, Frequency, and Q parameters.

To use the Dynamic EQ effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Filters, and double-click Dynamic EQ.
- 3 Set the desired options.

For more information, see [“Dynamic EQ options in Edit View” on page 121](#) or [“Dynamic EQ options in Multitrack View” on page 121](#).

Dynamic EQ options in Edit View

In Edit View, the Dynamic EQ effect provides the following options:

Frequency graph Sets the EQ frequency. The horizontal ruler (x-axis) represents the length of the selection, while the vertical ruler (y-axis) represents the frequency that's boosted or cut.

Gain graph Adjusts the amount of amplitude or attenuation the Dynamic EQ effect uses. The horizontal ruler (x-axis) represents the length of the selection, while the vertical ruler (y-axis) represents the dB level that's boosted or cut.

Q (Bandwidth) graph Adjusts the amount of Q used by the Dynamic EQ effect. The horizontal ruler (x-axis) represents the length of the selection, while the vertical ruler (y-axis) represents the frequency that's boosted or cut.

Flat Resets the graph to its default state.

Filter Type Lets you select from Low Pass, High Pass, and Band Pass:

- Low Pass preserves low frequencies and removes high frequencies.
- High Pass preserves high frequencies and removes low frequencies.
- Band Pass preserves a *band* (a range of frequencies) while attenuating all other frequencies. In the Stop Band text box, specify the number of decibels by which you want to attenuate other frequencies. For traditional band pass filtering, use the default setting, 30 dB.

Loop Graph When unselected, causes the graphs to be as long as the selection, so the dynamic changes to the equalization are graphed across the whole selection. If you select Loop Graph, the graphs are as long as you specify in the corresponding text box (which is enabled when Loop Graph is checked), so they loop repeatedly over the course of the selection.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Dynamic EQ options in Multitrack View

In Multitrack View, the Dynamic EQ effect provides the following options:


Gain Draws dynamic gain curves with yellow clip envelopes if Automated is selected. Lets you specify a static boost or cut if Automated is deselected.

Frequency Draws dynamic frequency curves with pink clip envelopes if Automated is selected. Lets you specify a static frequency if Automated is deselected.

Q (bandwidth) Draws dynamic bandwidth curves with purple clip envelopes if Automated is selected. Lets you specify a static bandwidth if Automated is deselected.

Filter Type Lets you select from Low Pass, High Pass, and Band Pass. For more information, see [“Dynamic EQ options in Edit View” on page 121](#).




To see clip envelopes for these options, click the Show FX Parameter Envelopes button  in the toolbar. For more information, see [“Automating mixes with clip envelopes” on page 207](#).

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the FFT Filter effect (Edit View only)

The graphic nature of the FFT (Fast Fourier Transform) Filter effect makes it easy to draw curves or notches that reject or boost specific frequencies. This effect can produce broad band-pass filters such as high- and low-pass filters (to maintain high and low frequencies, respectively), narrow band-pass filters (to simulate the sound of a telephone call), or notch filters (to eliminate very narrow frequency bands). The noise level of the FFT Filter effect is lower than that of 16-bit samples, so it introduces no noise when processing audio at 16-bit resolution or lower.



FFT Filter graph (De-Esser preset)



For optimal results, filter 32-bit samples. If the source audio is 8-bit or 16-bit, convert it to 32-bit first, and after you filter it, convert it back to 8-bit with dithering. You'll produce better results than processing at lower resolutions, especially if you perform more than one transform on the audio.

To use the FFT Filter effect:

- 1 In Edit View, select an audio range.
- 2 In the Effects tab of the Organizer window, expand Filters, and double-click FFT Filter.
- 3 Set the desired options.

For more information, see [“FFT Filter options” on page 122](#).

FFT Filter options

The FFT Filter effect provides the following options:

Passive, Logarithmic modes Measures frequency changes (boosts or cuts) in percentages (Passive) or dB (Logarithmic), where a setting of 100% or 0 dB represents no change.

View Initial Filter Graph, View Final Lets you set both an initial and a final filter setting if Lock To Constant Filter isn't selected. The rate at which the filter migrates from the initial to final settings depends on the Transition Curve settings.

Log Scale Displays the x-axis (frequency) in a logarithmic scale rather than a linear scale. A logarithmic scale more closely resembles the way the ear hears sound.

- To do finer editing in low frequencies, select Log Scale.
- For detailed, high-frequency work, or work with evenly spaced intervals in frequency, deselect Log Scale.

Flat Resets the graph to its default state.

Max, Min Sets the maximum and minimum values for the horizontal ruler (y-axis).

FFT Size Specifies the size of the FFT to use (represented as a power of two), thereby affecting processing speed and quality. For cleaner sounding filters, use higher values. Values between 1024 and 8192 work well.



Use a lower value (512 or so) for faster previews, and use a higher value after you achieve the settings you want for better quality when processing audio.

Windowing Function Determines the amount of transition width and ripple cancellation that occurs during filtering, with each one resulting in a different frequency response curve. These functions are listed in order from smallest width and greatest ripples to widest and least ripples.

The filters with the least ripples are also those that most precisely follow the drawn graph and have the steepest slopes, even though they are wider and pass more frequencies in a band-pass operation. The Hamming and Blackman filters give excellent overall results.

Lock To Constant Filter Applies a constant filter to the waveform. Deselect this option to set initial and final filter settings.

Morph Morphs the initial filter settings to the final filter settings. If this option is deselected, the settings simply change in linear fashion over time. For example, if Morph is unselected, and you have a spike at 10 kHz for the initial filter and a spike at 1 kHz for the final filter, the spike at 10 kHz will gradually decrease over time, and the spike at 1 kHz will gradually increase over time; frequencies between 1 kHz and 10 kHz won't be affected. If Morph is selected, however, the spike will gradually transition from 10 kHz down to 1 kHz, passing through the frequencies in between.



For a cool example of morphing, select Passive mode, and set an initial curve with the first half at 100% and the second half at 0%. For the final curve, set the right tenth or so at 100% with the rest at 0%. This combination selects high frequencies for the initial configuration, and low frequencies for the final configuration.

To get a nice blending from high to low, select Morph to include all the frequency combinations between the two filters. Click Transition Curve to view the actual settings that will be used over the duration of the selection.

Precision Factor Determines how accurately you want to filter over time when separate initial and final settings are used. A larger number (low factor) causes the filter settings to change roughly (or in chunks) from initial to final, while smaller numbers (higher factor) make the transition much smoother. In any case, the higher the precision factor, the longer the processing time, but the nicer the effect will sound.

Since the FFT function takes a large group of samples and filters them all at once, the precision factor determines how many samples from the entire group are actually saved. A factor of 2 means that 1/2 of the samples are saved, while a factor of 10 means that 1/10 of the samples are saved. Since you can have only one filter setting for the entire group of samples, use a more accurate (smaller) setting if the EQ curve varies wildly over short periods of time.

Transition Curve Opens the Transition Curve window, which displays a graphical representation of the transition from initial to final filter settings. The top graph shows time along the x-axis (where the left represents the start of the sample, and the right represents the end), and it shows where in the transition you are allowed the y-axis (where 0% represents the initial filter, and 100% represents the final filter). All points in between are a combination of initial and final filter arrangements. The readout below the graph displays the position of the cursor.

In the Transition Curve window, select from the following options:

- Flat resets the curve to its default state.
- Graph Response At Point causes the bottom graph to change in response to the cursor position in the top graph and shows the filter at any given point in the transition. Depending on the position you select, you can specify a morphing transition or a linear transition.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Graphic Equalizer effect

The Graphic Equalizer effect boosts or cuts specific frequency bands and provides a visual representation of the resulting EQ curve. Unlike the Parametric Equalizer, the Graphic Equalizer uses preset frequency bands for quick and easy equalization. The fixed Q settings ensure that no drop outs exist at intermediate frequencies. You can space frequency bands at intervals of one octave, one-half octave, or one-third octave.

The Graphic Equalizer effect is an FIR (Finite Impulse Response) filter, which maintains phase accuracy—unlike an IIR (Infinite Impulse Response) filter, which can introduce phase errors, adding a ringing quality to audio.

To use the Graphic Equalizer effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Filters, and double-click Graphic Equalizer.
- 3 Set the desired options.

For more information, see [“Graphic Equalizer options” on page 124](#).

Graphic Equalizer options

The Graphic Equalizer effect provides the following options:

Bands tabs Let you access the number of EQ bands you need. Graphic Equalizer bands are spaced at intervals of one octave (10 bands), one-half octave (20 bands), or one-third octave (30 bands). The 10 Bands setting offers more general equalizing, while the 20 and 30 Band settings let you zoom in more precisely on specific frequency ranges. Other than the bands settings, the controls in each tab of the Graphic Equalizer window are identical.

Reset All To Zero Sets all sliders to 0 dB so no equalization occurs.

Band Determines the band to be modified with the Gain value.

Gain Sets the exact value for the gain (measured in decibels) by which to modify the chosen band.

Graph Actual Response Calculates the actual response of the equalizer rendered and displays it in the window above the band sliders. Because the Graphic Equalizer effect is an FIR filter, the response may not actually match the desired equalization curve at lower accuracy levels.

Accuracy Sets the accuracy level for equalization. Higher accuracy levels (longer FIR filters) give better frequency response in the lower ranges, but they require more processing time. If you equalize only higher frequencies, you can use lower accuracy levels.



If you equalize extremely low frequencies, set Accuracy to between 500 and 5000 points.

Range Defines the range of the slider controls. Enter any value between 4 and 180 dB. (By comparison, standard hardware equalizers have a range of about 30 to 48 dB.)

Master Gain Compensates for an overall volume level that is too soft or too loud after the EQ settings are adjusted. The default value of 0 dB represents no master gain adjustment.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Notch Filter effect

The Notch Filter effect removes up to six user-defined frequency bands, in addition to standard telephone DTMF tones. Use this effect to remove very narrow frequency bands, such as a 60 Hz hum, while leaving all surrounding frequencies untouched.

To use the Notch Filter effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Filters, and double-click Notch Filter.
- 3 Set the desired options.

For more information, see [“Notch Filter options” on page 125](#).

Notch Filter options

The Notch Filter effect provides the following options:

Fix Attenuations To Determines if notches have equal or individual attenuation levels.

Frequency Specifies the center frequency for each notch.

Attenuation Specifies the amplitude reduction for each notch.

Notch Width Determines frequency range for all notches. The three options range from Narrow for a second order filter, which removes some adjacent frequencies, to Super Narrow for a sixth order filter, which is very specific.



Use no more than 30 dB attenuations for a Narrow setting, no more than 60 dB for Very Narrow, and no more than 90 dB for Super Narrow. Greater attenuation can remove a wide range of neighboring frequencies.


DTMF Lower Tones, DTMF Upper Tones Filters the standard lower and upper DTMF telephone tones. These options are useful if you prepare audio for radio.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Parametric Equalizer effect

The Parametric Equalizer provides maximum control over tonal equalization. Unlike the Graphic Equalizer, which provides a fixed number of frequencies and Q bandwidths, the Parametric Equalizer gives you total control over frequency, Q, and gain settings. For example, you can simultaneously reduce a small range of frequencies centered around 1000 Hz, boost a broad low-frequency shelf centered around 80 Hz, and insert a 60 Hz notch filter.

The Parametric Equalizer uses second-order IIR filters, which are very fast and provide very precise resolution, even at lower frequencies. For example, you can precisely boost a range of 40 to 45 Hz.



Parametric EQ graph (Old Time Radio preset)

To use the Parametric Equalizer effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Filters, and double-click Parametric Equalizer.
- 3 Set the desired options.

For more information, see [“Parametric Equalizer options” on page 126](#).

Parametric Equalizer options

The Parametric Equalizer effect provides the following options:

Graph Depicts frequency along the horizontal ruler (x-axis) and amplitude along the vertical ruler (y-axis), with the curve representing the amplitude change at specific frequencies. Frequencies in the graph range from lowest to highest in a logarithmic fashion (evenly spaced by octaves).

Low Shelf Cutoff (Determines the Low Shelf amplitude in decibels. Increase or decrease lows (bass) by adjusting the slider at the left of the graph, or type a value in the text box below the slider. Use the Low Shelf Cutoff to reduce hiss, amplifier noise, and the like.

High Shelf Cutoff Determines the High Shelf amplitude in decibels. Increase or decrease the highs (treble) by adjusting the slider at the right of the graph, or type a value in the text box below the slider. Use the High Shelf Cutoff to reduce low-end rumble, hum, or other unwanted low-frequency sounds.

Center Frequency Places up to five intermediate bands into the EQ circuit, giving you very fine control over the shape of the equalization curve. Select the box next to a slider to activate the band and its corresponding volume slider and Hz text box, both of which control the center frequency at which the boost or cut occurs. The vertical sliders in the upper right of the Parametric Equalizer dialog box control the amount of boost or cut. You can also specify the amount of boost or cut amount in the text box below each slider.

Width Controls the width of the affected frequency band, measured in either Q or Width values. Low Q values (or a high Width values) affect a larger range of frequencies. Very high Q values (above 100) affect only a very narrow band and are ideal for notch filters where only a particular frequency needs to be removed, like a 60 Hz hum.



Be aware that whenever a very narrow band is boosted, it tends to ring or resonate at the audio at that frequency. Q values of 1 to 10 are used most often for general equalization.

Constant Width, Constant Q Describes a frequency band's width as either a Q value (which is a ratio of width to center frequency) or an absolute width value in Hz. Constant Q is the most common setting, but you may want to use Constant Width if, for example, you want the length of ringing to be a constant, no matter what frequency is boosted.

Master Gain Compensates for an overall volume level that might be too loud or too soft after you adjust the EQ settings.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Quick Filter effect (Edit View only)

The Quick Filter is an 8-band graphic equalizer that you can easily customize to suit many filtering needs. Unlike a standard graphic equalizer, settings for the individual frequency bands interact with nearby frequencies. For example, significantly boosting the level of the highest 22 kHz frequency band moderately boosts the level of lower frequencies. This behavior helps you to quickly and easily enhance audio tone.



To change the equalization setting over time, use different Initial and Final settings. Using this approach, you can create many interesting effects, such as bass-heavy equalization that gradually changes to flat equalization at the introduction of a song.

To use the Quick Filter effect:

- 1 In Edit View, select an audio range.
- 2 In the Effects tab of the Organizer window, expand Filters, and double-click Quick Filter.
- 3 Set the desired options.

For more information, see [“Quick Filter options” on page 127](#).

Quick Filter options

The Quick Filter effect provides the following options:

Lock To These Settings Only Determines whether settings are equalized or varied.

- If this option is selected, the entire selected range is equalized at the settings shown.
- If this option is unselected, you can specify separate initial and final equalization settings, so that the selection can smoothly glide from the initial equalization setting to the final setting over the selected range. Click the Initial Settings and Final Settings tabs to specify initial and final settings.

Initial Settings tab Appears if Lock To These Settings Only isn't selected. Click this tab to specify the initial EQ settings.

Final Settings tab • Appears if Lock To These Settings Only isn't selected. Click this tab to specify the final EQ settings.

Equalizer band sliders Increase or decrease the frequency specified beneath each slider. Amplitude appears above each slider.

Master Gain sliders Adjust the equalizer's overall level for both the left and right channels of stereo waveforms.



If you increase the EQ frequencies of a waveform, the waveform's volume usually increases too, potentially leading to clipping. Use the Master Gain sliders to reduce the level before applying the effect.

Lock L/R Lets you adjust channels together or separately. Select this option to adjust the channels together, maintaining the same settings for each. Leave this option unselected to adjust each channel separately.




To produce a low-pass filter, set the higher frequency scroll bars to -30. Similarly, you can create a high-pass filter by reducing the lower frequencies.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Scientific Filters effect (Edit View only)

The Scientific Filters effect provides high-order IIR (Infinite Impulse Response) filters for precise band-pass, notch, or high- or low-pass filtering. The most common types of high-order filters are available: Bustle, Butterscotch, Chebychev 1, and Chebychev 2. Each type has different characteristics for filter attenuation and the steepness of transition bands at cutoff points. Butterworth usually provides the best compromise between quality and precision.

On the Scientific Filters graph, one line shows frequency response (measured in decibels), and the other line shows either phase (measured in degrees) or group delay (measured in milliseconds), depending on whether the Phase or Delay option is selected. Increase the graph's display range by selecting Extended Range.



Scientific Filters graph for Butterworth filter (Remove Subsonic Rumble preset)

A. Group Delay (milliseconds) **B.** Frequency Response (dB)

To use the Scientific Filters effect:

- 1 In Edit View, select an audio range.
- 2 In the Effects tab of the Organizer window, expand Filters, and double-click Scientific Filters.
- 3 Set the desired options.

For more information, see [“Scientific Filters options” on page 129](#).

Scientific Filters options

The Scientific Filters effect provides the following options:

Bessel Provides accurate phase response with no ringing or overshoot. However, the pass band slopes at its edges, where rejection of the stop band is poorest of all filter types. These qualities make Bessel a good choice for percussive, pulse-like signals. For other filtering tasks, use Butterworth.

Butterworth Provides a flat pass band with minimal phase shift, ringing, and overshoot. This filter type also rejects the stop band much better than Bessel and only slightly worse than Chebychev 1 or 2. These overall qualities make Butterworth the best choice for most filtering tasks.

Chebychev 1 Provides the best stop band rejection but the worst phase response, ringing, and overshoot in the pass band. Use this filter type only if rejecting the stop band is more important than maintaining an accurate pass band.

Chebychev 2 Combines a Butterworth filter in the pass band with notch filters in the stop band. In between the notches of the stop band, some phase-shifted signal remains, but at highly attenuated levels.

Low Pass Passes the low frequencies and removes high frequencies. You must specify the cutoff point at which the frequencies are removed.

High Pass Passes high frequencies and removes low frequencies. You must specify the cutoff point at which the frequencies are removed.

Band Pass Preserves a *band*, a range of frequencies, while removing all other frequencies. You must specify two cutoff points to define the edges of the band.

Band Stop Rejects any frequencies within the specified range. Also known as a notch filter, Band Stop is the opposite of Band Pass. You must specify two cutoff points to define the edges of the band.

Cutoff Defines the frequency that serves as a border between passed and removed frequencies. At this point the filter switches from passing to attenuating, or vice versa. In filters requiring a range (Band Pass and Band Stop), Cutoff defines the low frequency border, while High Cutoff defines the high frequency border.

High Cutoff Defines the high frequency border in filters that require a range (Band Pass and Band Stop).

Order Determines the filter's precision. The higher the order, the more precise the filter (with steeper slopes at the cutoff points, and so on). However, very high orders can also have high levels of phase distortion.

Transition Bandwidth (Butterworth and Chebychev only) Sets the width of the transition band. (Lower values have steeper slopes.) If you specify a transition bandwidth, the Order setting is filled in automatically, and vice-versa. In filters that require a range (Band Pass and Band Stop), this serves as the lower frequency transition, while High Width defines the higher frequency transition.

High Width (Butterworth and Chebychev only) In filters that require a range (Band Pass and Band Stop), this option serves as the higher frequency transition, while Transition Bandwidth defines the lower frequency transition.

Pass Ripple/Actual Ripple (Chebychev only) Determines the maximum allowable amount of ripple. Ripple is the effect of unwanted boosting and cutting of frequencies near the cutoff point.

Stop Attn (Butterworth and Chebychev only) Determines how much gain reduction to use when frequencies are removed.

Master Gain Compensates for an overall volume level that might be too loud or too soft after you adjust the filter settings.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Optimizing amplitude

Amplitude effects let you optimize audio volume for specific mediums such as radio and CD, produce detailed fade outs, and more.

Using the Amplify/Fade effect (Edit View only)

The Amplify/Fade effect produces either constant amplification changes (such as fixed boosts) or precise fades.



Though the Amplify/Fade effect isn't available in Multitrack View, you can use real-time envelopes to accomplish the same task. (See [“Automating mixes with clip envelopes” on page 207](#).)

To use the Amplify/Fade effect:

- 1 In Edit View, select an audio range.
- 2 In the Effects tab of the Organizer window, expand Amplitude, and double-click Amplify/Fade.
- 3 Set the desired options.

For more information, see [“Amplify/Fade options” on page 130](#).

Amplify/Fade options

The Amplify/Fade effect provides the following options:

Amplification (Constant Amplification tab) Determines the amount of volume change to apply to the selection. An amplification value greater than 100% or 0 dB increases the volume; a value below 100% or 0 dB reduces the volume.

Initial Amplification (Fade tab) Determines the amount of volume change to apply to the beginning of the selection. An amplification value greater than 100% or 0 dB increases the volume; a value below 100% or 0 dB reduces the volume.



To fade in audio, specify an Initial Amplification value that's lower than the Final Amplification value. To fade out, do the opposite.

Final Amplification (Fade tab) Determines the amount of volume change to apply to the end of a selection.

Linear Fades (Fade tab) Causes the waveform's sample values to be faded in an even, linear fashion, producing a smooth slope from beginning to end.

Logarithmic Fades (Fade tab) Applies a logarithmic-style fade (also known as a “power fade”). If you select this option, the amplitude of a signal fades at a constant rate, producing a steeper slope at one end of the fade (depending on whether you fade in or out).

DC Bias Adjust Ensures that new recordings are perfectly centered. Some recording hardware may introduce a DC bias, causing the recorded waveform to appear to be above or below the normal center line in the wave display. Many waveform transformations require that the signal be centered.

- Select Absolute and then specify the final DC percentage in the L and R boxes. This option lets you cancel out DC that’s not constant throughout a waveform. An extreme low-cut filter achieves this result. Keep in mind that the actual amount adjusted varies with each sample. For example, if you have a significant DC change in one area of the wave, at that boundary where the DC changes, the Absolute option makes all parts the same. (However, there will be a dip or peak right at the boundary point.) To introduce a DC bias by skewing the entire selected waveform above or below the center line, enter a positive or negative percentage. For example, a setting of 50% moves the entire waveform up halfway, and one of –50% moves it down halfway.
- Select Differential and click Find Zero Now to analyze the entire selected area to get the DC offset, and adjust every sample by the inverse of that exact amount. The correct L and R percentages are entered automatically.

Peak Level Sets the peak level used for normalizing audio.

Calculate Now Scans the selection and adjusts the amplification sliders to normalize the selection according to the peak level.

Lock Left/Right Lets you adjust channels together or separately. Select this option to adjust the channels together, maintaining the same settings for each. Leave this option unselected to adjust each channel separately. Separate adjustments let you tweak the stereo balance or create cool panning effects.

View All Settings In dB Causes amplification values to appear in decibels; otherwise, they appear as a percentage of the original waveform.


Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Envelope effect (Edit View only)

The Envelope effect lets you precisely control amplitude over time, enabling you to combine a wide range of amplification effects, such as multiple fades and boosts. The top of the Envelope graph represents 100% (normal) amplification; the bottom represents 100% attenuation (silence).



Though the Envelope effect isn't available in Multitrack View, you can use real-time track envelopes to accomplish the same task. (See ["Automating mixes with clip envelopes" on page 207.](#))



Envelope graph (Bell Curve preset)

To use the Envelope effect:

- 1 In Edit View, select an audio range.
- 2 In Effects tab of the Organizer window, expand Amplitude, and double-click Envelope.
- 3 Set the desired options.

For more information, see ["Envelope options" on page 132.](#)

Envelope options

The Envelope effect provides the following options:

Envelope graph Depicts time along the horizontal ruler (x-axis) and the new output level along the vertical ruler (y-axis), with the blue line representing amplitude change.

Flat Resets the graph to its default state.

Amplification Specifies where the top of the graph is, measured in percentage. When you draw an envelope curve, the top of the graph is whatever percentage you specify and the bottom is 0%. The default value of 100% means no change in volume occurs.



Use the Envelope effect to make tones generated with Adobe Audition sound more realistic.


Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see ["Working with effects" on page 33.](#)

Using the Dynamics Processing effect

The Dynamics Processing effect varies the output level of a waveform based on its input level. You can use this effect to limit or compress dynamic range, producing a consistent level of perceived loudness. You can also expand or gate the signal so that low-level signals are reduced in level, increasing perceived dynamic range, or eliminating signals with noise that falls below a specific threshold.



The Dynamics Processing effect can produce subtle changes that you notice only after repeated listening. When applying this effect in Edit View, use a copy of the original file so you can return to the original audio if necessary.



Dynamics Processing graph (Classic SoftKnee preset)

To use the Dynamics Processing effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Amplitude, and double-click Dynamics Processing.
- 3 Set the desired options.

For more information, see [“Dynamics Processing options” on page 133](#).

Dynamics Processing options

The Dynamics Processing effect provides the following options:

Graphic tab Contains a graph that depicts input level along the horizontal ruler (x-axis) and the new output level along the vertical ruler (y-axis). The default graph, with a line from the lower left to the upper right, depicts a signal that has been left untouched, since every input value goes to the exact matching output value. Adjusting the graph adjusts the input or output assignments, thereby altering the dynamic range.

For example, you can boost all input that has a level of around –20 dB, leaving everything else unchanged. You can also draw an inverse line (a line from the upper left to the lower right) that will dramatically boost low amplitudes while dramatically suppressing high amplitudes (that is, all quiet sounds will be loud, and all loud sounds will be quiet).

Flat Resets the graph to its default state.

Invert Inverts the graph.

Note: You can invert a graph only if it has points in the two default corners (–100, –100 and 0, 0) and if its output level increases from left to right (that is, each control point must be higher than the one to its left).

Traditional tab Lets you specify ratios and thresholds. You can choose Compress, Flat, or Expand for up to six sections or stages, each with its own ratio and threshold setting. For example, to create a 3:1 compressor above –20 dB, choose Compress, and then specify a 3:1 ratio and a threshold of –20 dB. If you also want to expand 2:1 below –20 dB, choose Expand, and then specify a 2:1 ratio.

Note: Threshold settings must decrease as you move down the list.

Attack/Release tab Provides settings for Gain Processor and Level Detector.

Gain Processor Amplifies or attenuates the signal depending on the amplitude detected. Specify the following settings as desired:

- Output Gain is gain applied to the output signal. It's the last step performed on the audio.
- Attack Time is applied just before output and determines the time in milliseconds that it takes for the processed output signal to reach its specified output volume. If a quiet portion suddenly drops 30 dB, the specified amount of time passes before the output actually drops to its corresponding volume level.
- Release Time is applied just before output and determines how long it takes for the end of a previous output level to reach the specified output volume. For example, where the Attack Time is the time it takes for the start of a pulse to reach the desired output volume, the Release Time is the time it takes for the end of the pulse to reach the desired level.

Note: If the sum of Attack and Release times is too short (less than about 30 milliseconds), audible effects, such as a vibrating noise, can be heard at around 1000 Hz/milliseconds total. For example, if the Attack and Release times are each set to 5 milliseconds (for a sum of 10 milliseconds), then a vibrating noise can be heard at 100 Hz.

- Joint Channels: Uses both channels to find a single input dB value, so that both channels are amplified together by the same amount (thus preserving the stereo center-channel image). For instance, a loud drum beat on the left channel will cause the right channel to be reduced in level by an equal amount.

Level Detector Determines the current waveform input amplitude, which is used as the input side of the dynamics processor. Specify the following settings as desired:


- Input Gain is gain applied to the signal before it goes into the Level Detector (the section that detects the current level).
- Attack Time is applied when the current amplitude information is retrieved and determines the time in milliseconds that it takes for the processed output signal to reach its specified output volume. If a quiet portion suddenly drops 30 dB, the specified amount of time passes before the output actually drops to its corresponding volume level.
- Release Time is applied when the current amplitude information is retrieved and determines how long it takes for the end of a previous output level to reach the specified output volume. For example, where the Attack is the time it takes for the start of a pulse to reach the desired output volume, the Release is the time it takes for the end of the pulse to reach the desired level.

- Peak mode is provided for backward compatibility. It is a graph interpretation method that is slightly outdated and a bit more difficult to use than RMS. It is equivalent to twice the RMS value (for example, –20 dB in RMS mode equals –40 dB in Peak mode).
- RMS is a graph interpretation method that more closely matches the way people hear volume. This mode causes the output to be exactly the RMS amplitude that is specified in the graph. For example, a limiter (flat horizontal line) at –10 dB causes the RMS amplitude of the result to average –10 dB (where 0 dB is a maximum amplitude sine wave without clipping).


Lookahead Time Helps to handle sharp spikes that might occur at the onset of a louder signal and go beyond the limits of the compressor settings. While this approach may be desirable to enhance the impact of drum hits, for example, it isn't desirable if you use limiting to reduce the maximum amplitude of audio.

Note: *The spikes occur because it takes a little time to determine and react to the current signal level (that is, to determine the Level Detector's attack value and the Gain Processor's attack value). Lookahead Time actually causes the attacks to start before the audio gets loud instead of right on top of the transient. Otherwise, with a Lookahead Time of 0, a spike stays loud until all of the attack times have elapsed.*


Band Limiting tab Lets you limit dynamics manipulation to a specified range:

 In Edit View, you can apply multiband compression by running a script. When creating the script, include multiple passes of the Dynamics Processing effect, each with different Low and High Cutoff settings. (See [“Working with scripts” on page 254.](#))

- Low Cutoff is the lowest frequency that dynamics processing will affect. You can define a *band*, or range, to which compression or expansion is applied within the current frequency range.
- High Cutoff is the highest frequency that dynamics processing will affect. To use the entire frequency range of the source material, enter a value of 0.

 To use the entire frequency range of the source material, set the Low Cutoff to 0 and the High Cutoff to 1/2 the current sample rate (24,000 for 48 kHz, 11,025 for 22 kHz, and so on).

Create Envelope Only Applies any dynamics processing and returns the result as an amplitude envelope. To modulate this envelope with another sound's amplitude, copy it, choose Edit > Paste, and select Modulate. If you click Preview with Create Envelope Only selected, you'll hear pink noise instead of the audio that's highlighted in the Wave Display.

 Use the RadioLimit, Fast Release, Boost preset to simulate the processed sound of a contemporary FM radio station.

Note: *For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33.](#)*

Using the Hard Limiting effect

The Hard Limiting effect drastically attenuates audio that rises above a defined threshold, leaving audio below the threshold unaffected. This effect is particularly useful for increasing perceived volume because you can amplify audio beyond the digital maximum, 0 dBFS, and you can lower areas that would otherwise be clipped. For example, when you convert from 32-bit to 16-bit audio, particularly loud 32-bit passages can cause 16-bit clipping. To prevent clipping, you can either use the Normalize effect to reduce the amplitude of the entire file (lowering perceived volume), or you can use the Hard Limiting effect to reduce amplitude only for loud passages (increasing perceived volume).

To use the Hard Limiting effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Amplitude, and double-click Hard Limiting.
- 3 Set the desired options.

For more information, see [“Hard Limiting options” on page 136](#).

Hard Limiting options

The Hard Limiting effect provides the following options:

Limit Max Amplitude To Sets the maximum sample amplitude allowed.



To avoid clipping when working with 16-bit audio, set this value to no more than –0.1 dB; if you set it to –0.5 dB, you’ll have a little more clearance for any future edits.

Boost Input By Preamplifies audio before you limit it, so you can make a selection louder and ensure that no clipping occurs (similar to what is commonly done with the audio on TV commercials).

Lookahead Time Sets the amount of time (measured in milliseconds) generally needed to attenuate the audio before the loudest peak is hit.

Note: Make sure that the value is at least 5 milliseconds. If this value is too small, audible distortion effects may occur.

Release Time Sets the time (measured in milliseconds) needed for the attenuation to rebound back 12 dB (or roughly the time needed for audio to resume normal volume if an extremely loud peak is encountered). In general, a setting of around 100 (the default) works well and preserves very low bass frequencies.

Note: If this value is too large, audio may remain very quiet and not resume normal levels for a while.

Link Left And Right Links the loudness of both channels together, preserving the stereo image.

Gather Statistics Now Updates the Clipping Statistics values, which indicate what percentage of the audio would clip if limiting weren’t performed. Click this option after you change any of the input parameters.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Normalize effect (Edit View only)

The Normalize effect lets you set a peak level for a file or selection. When you normalize audio to 100%, you achieve the maximum amplitude that digital audio allows—0 dBFS.

The Normalize effect amplifies the entire file or selection equally. For example, if the original audio reaches a loud peak of 80% and a quiet low of 20%, normalizing to 100% amplifies the loud peak to 100% and the quiet low to 40%.



To apply RMS normalization, you must use the Group Waveform Normalize command. If desired, you can apply that command to only one file. (See [“Normalizing groups of files” on page 250.](#))

To use the Normalize effect:

- 1 In Edit View, select an audio range.
- 2 In the Effects tab of the Organizer window, expand Amplitude, and double-click Normalize.
- 3 Set the desired options.

For more information, see [“Normalize options” on page 137.](#)

Normalize options

The Normalize effect provides the following options:

Normalize To Sets the percentage by which you want to normalize. For example, specify 50% to amplify a selection by no more than 50% of maximum (resulting in a 3 dB attenuation from maximum output). Specify 100% (the default) to apply the most amount of amplification possible without clipping.

Decibels Format Displays the Normalize value in decibels instead of as a percentage.

Normalize L/R Equally Causes both channels of a stereo waveform to be used when the amplification amount is calculated. If this option is unselected, the amount is calculated separately for the channels, potentially amplifying one considerably more than the other.

DC Bias Adjust Lets you adjust the position of the waveform in the wave display. Some recording hardware may introduce a DC bias, causing the recorded waveform to appear to be above or below the normal center line in the wave display. To center the waveform, set the percentage to zero. To skew the entire selected waveform above or below the center line, specify a positive or negative percentage.



If you plan to use normalized audio on a CD, you might want to normalize the waveforms to no more than 96% as some audio compact disc players inaccurately reproduce bits that are processed to 100% (maximum) amplitude.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33.](#)

Applying Stereo, Pitch, and Delay Effects

About using stereo, pitch, and delay effects

Adobe Audition contains a wide variety of effects that let you change stereo imagery, adjust pitch, and add delay (for example, reverb and echo). Dialog boxes for these effects share many common options, such as graphs, spline curves, presets, and previews. For information on these shared options, see [“Working with effects” on page 33](#).

Note: You can apply effects differently in Edit View and Multitrack View, and some effects dialog boxes have different options in each view. For information on applying effects in Edit View, see [“Selecting audio data” on page 76](#). For information on applying effects in Multitrack View, see [“Using real-time effects” on page 205](#).

Changing stereo imagery

Adobe Audition lets you change the apparent location, or *stereo imagery*, of sounds coming from the speakers. For instance, you can move a sound from the center to the left or right speaker or even make sounds seem to circle a listener’s head.

Note that all of the stereo imagery effects except the Doppler Shifter effect work only on stereo files.

Using the Binaural Auto-Panner effect (Edit View only)

The Binaural Auto-Panner effect spatially lets you designate sound spatially on the left and the right in a seemingly circular pattern over time. In order to spatially encode the sound, either the left or right channel is delayed so that the sounds reach each ear at different times, tricking the brain into thinking they are coming from either side.

To use the Binaural Auto-Panner effect:

- 1 In Edit View, select a stereo range.
- 2 In the Effects tab of the Organizer window, expand Amplitude, and double-click Binaural Auto-Panner.
- 3 Set the desired options.

For more information, see [“Binaural Auto-Panner options” on page 138](#).

Binaural Auto-Panner options

The Binaural Auto-Panner effect provides the following options:

Frequency graph Represents time along the x-axis (horizontal edge) and frequency along the y-axis (vertical edge). Specify the highest and lowest frequencies represented on the graph with the options under Bottom Graph Settings and Top Graph Settings. The readout below the graph displays the current x, y position of the pointer.



Do Delay Only Performs only a delay on the audio.

Flat Resets the graph to its default state.

Bottom Graph Settings, Top Graph Settings Control the low and high binaural frequencies. Bottom Graph Settings correspond to the lower part of the graph, and Top Graph Settings correspond to the top:

- Pan Cycling Rate represents how often the sound moves from one channel to the other and back again. For example, a Pan Cycling Rate of 1 Hz moves a mono source from left to right and back in 1 second, while a Pan Cycling Rate of 0.1 Hz pans it in 10 seconds.
- Intensity controls the degree of the binaural encoding. Higher intensities work well with lower binaural frequencies.
- Centering tricks the brain into thinking the signal is coming from the left or right.



To create an interesting effect, mix a file that is binaurally processed to the left with one that is binaurally processed to the right (where the Pan Cycling Rate for each is within 2 Hz of each other).

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Channel Mixer effect

The Channel Mixer effect alters the left and right balance of a stereo waveform, letting you create new stereo mixes by using the existing right and left channels as input sources. By recombining and inverting the channels, you can create some interesting stereo-imaging effects.

To use the Channel Mixer effect:

- 1 Select a stereo range.
- 2 In the Effects tab of the Organizer window, expand Amplitude, and double-click Channel Mixer.
- 3 Set the desired options.

For more information, see [“Channel Mixer options” on page 139](#).

Channel Mixer options

The Channel Mixer effect provides the following options:

New Left Channel Determines the percentage of the current left and right channels to mix into the new left channel. For example, an L value of 50 and an R value of 50 makes the new left channel contain equal audio from both the current L and R channels. In contrast, an L value of 0 and a R value of 100 makes the new left channel contain only the audio of the current right channel.

New Right Channel Determines the percentage of the current left and right channels to mix into the new right channel.

Invert Inverts a channel's phase polarity (that is, the peaks become valleys and the valleys become peaks). Inverting both channels causes no perceived difference in sound. Inverting only one channel, however, places the channels out of phase and greatly changes the sound.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Pan/Expand effect

The Pan/Expand effect lets you shift the center channel of a stereo waveform. It also lets you expand or narrow the stereo separation of the left and right channels.

Center channel panning uses the *surround* and *center* channels of a stereo recording, where the *surround* channel is the difference of the two original channels, and the *center* channel is the sum of them. You can think of a stereo recording as having four channels (left, right, center, and surround), and this effect lets you pan these channels around. For example, pan hard left to get the original center channel to come out the left speaker and the original surround channel to come out the right. This type of panning can provide added realism to original stereo recordings.

Expanding works by subtracting or adding differing amounts of right and left channel signals, so sounds occurring on the right or left are cut or boosted. You can alter both of these elements dynamically over time by using the respective graph.

To use the Pan/Expand effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Amplitude, and double-click Pan/Expand.
- 3 Set the desired options.

For more information, see [“Pan/Expand options” on page 140](#).

Pan/Expand options

The Pan/Expand effect provides the following options:

Center Channel Pan graph (Edit View only) Represents the pan position of the center channel of a stereo waveform over time. The graph's x-axis (horizontal) represents the length of the waveform or selection, while the y-axis (vertical) represents the percentage of the pan from center. You can use the graph to position the center channel anywhere from hard left (–100%) to hard right (100%), with the corresponding surround channel moving right to left in the opposite direction. Use this method for panning original stereo data more realistically than amplitude panning allows.

Stereo Expand graph (Edit View only) Shows the expand level over time and amplifies (>100%) or removes (<100%) the differences between channels. The graph's x-axis (horizontal) represents the length of the waveform or selection, while the y-axis (vertical) represents the percentage of stereo expansion. With some material, you can create a stereo expanding effect by increasing the differences between the left and right channels. The expansion level can vary over time for interesting effects (growing from a mono signal to a very wide stereo signal, for example).

Flat (Edit View only) Resets the graph to its default state.

Automated Supports parameter automation when you use Pan/Expand as a real-time effect in Multitrack View. Select Automated in the Center Channel Pan section of the Multitrack View version of the Pan/Expand dialog box to draw the Center Channel Pan curve using a yellow envelope on waveform clips. Select it in the Stereo Expand section of the Multitrack View version of the Pan/Expand dialog box to draw the Stereo Expand curve using a pink envelope on waveform clips. Note that you must have View > Show FX Parameter Envelopes in Multitrack View turned on to use parameter automation.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Stereo Field Rotate effect

This effect lets you rotate the stereo field of an audio file. The stereo field denotes where in space instruments or other sources are placed within the left and right images of a stereo waveform. By manipulating the Rotation graph, you can affect how the instruments seem to move over time.

To use the Stereo Field Rotate effect:

- 1 Select a stereo range.
- 2 In the Effects tab of the Organizer window, expand Amplitude, and double-click Stereo Field Rotate.
- 3 Set the desired options.

For more information, see [“Stereo Field Rotate options” on page 141](#).


Stereo Field Rotate options

The Stereo Field Rotate effect provides the following options:

Rotation graph (Edit View only) Represents the rotation of the stereo field over time. Use the graph to position the stereo field anywhere from hard left (the top of the graph) to hard right (the bottom of the graph) at any point in time. The x-axis (horizontal border) of the graph shows the waveform’s timeline, while the y-axis (vertical border) displays the number of degrees off center for the left and right channels.

Note: For more information about graph controls (such as how to add and remove control points), see the *“Looking at the Work Area”* chapter.

Invert Left/Right Reverses the graph so that an upward line spins the stereo field clockwise instead of counter-clockwise.

 If the Phase Analysis window (Analyze > Show Phase Analysis) is open, you can watch the stereo field rotate. So, if a graph is drawn to rotate more and more to the right, the phase rotates clockwise (to the right). But, if you selected Invert Left/Right, the phase rotates counter-clockwise (to the left). However, you’ll actually hear the audio start panning to the right if the graph begins at 0.

Flat Resets the graph to its default state.

Range Sets the range of the y-axis from 45 degrees to 360 degrees.

Note: If you specify a stereo field rotation of 180 degrees, both the left and right channels are 180 degrees out of phase, inverting the waveform. (This result is the same as choosing the Invert effect from the Effects menu). At 90 degrees right, only the right channel is inverted. At 90 degrees left, just the left channel is inverted. A range of 45 degrees produces

results identical to the Pan/Expand effect, panning the center channel left or right while panning the surround channel the opposite direction.

Loop Graph Makes the graph only as long as specified so it loops repeatedly over the course of the selection. Selecting this option also enables the option for specifying Period (how long the graph should be), Frequency (how fast the cycle should move), Tempo (how many beats per minute the loop should have), or Total Cycles (how many times you want the graph to loop in the given selection). If Loop Graph is unselected, the graph is the same length as the selection, so dynamic changes to the delay are graphed across the whole selection.

Automated (Multitrack View only) Supports parameter automation if you use Stereo Field Rotate as a real-time effect in Multitrack View. Select Automated to define the Center Channel Pan curve by dragging a yellow envelope line on waveform clips. If this option is unselected, the dialog box displays a slider for setting the stereo field instead of the other options described. (You must have View > Show FX Parameter Envelopes In Multitrack View enabled to access the envelope in the dialog box.)

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Center Channel Extractor effect

The Center Channel Extractor effect keeps or removes frequencies that are common to both the left and right channels—in other words, sounds that are panned center. Often voice, bass, and lead instruments are recorded this way. As a result, you can use this effect to bring up the volume of the vocals, lead bass, or kick drum or remove any of them from the stereo mix.

To use the Center Channel Extractor effect:

- 1 Select a stereo range.
- 2 In the Effects tab of the Organizer window, expand Filters, and double-click Center Channel Extractor.
- 3 Set the desired options.

For more information, see [“Center Channel Extractor options” on page 142](#).

Center Channel Extractor options

The Center Channel Extractor effect provides the following options:

Get Audio Phased At Specifies the phase degree, pan percentage, and delay time for the audio you want to extract or remove. Set this option to Center (zero degrees) to work with audio that is panned to the exact center. To extract surround audio from a matrix mix, set this option to Surround (180 degrees) to work with audio that is exactly out of phase between the left and right channels. Set this option to Custom to modify phase degree and pan percentage, which can range from –100% (hard left) to 100% (hard right).

Frequency Range Sets the range you want to extract or remove. Predefined ranges include Male Voice, Female Voice, Bass, and Full Spectrum, and Custom. Set this option to Custom to define a frequency range.

Center Channel Level Specifies how much of the selected signal you want to extract or remove. Move the slider to the left (negative values) to remove center channel frequencies and to the right (positive values) to remove panned stereo material.

Volume Boost Mode Boosts center channel material if the Center Channel Level slider is set to a positive value and boosts panned stereo material if the slider is set to a negative values. This option is especially useful for boosting vocals.

Crossover Controls how much bleed through to allow. Move the slider to the left to increase audio bleed through and make it sound less artificial. Move the slider to the right to further separate center channel material from the mix.

Phase Discrimination In general, higher numbers work better for extracting the center channel, whereas lower values work better for removing the center channel. Lower values allow more bleed through and may not effectively separate vocals from a mix, but they may be more effective at capturing all the center material. In general, a range from 2 to 7 works well.

Spectral Decay Rate Keep at 0% for faster processing and to take advantage of multiple CPU and hyperthreaded computers. Set the value between 80% and 98% to help smooth out background distortions.

Amplitude Discrimination And Amplitude Band Width Sums the left and right channels, and creates a 180 degree-out-of-phase third channel that Audition uses to remove similar frequencies. If the volume at each frequency is similar, audio in common between both channels is also considered. Lower values for Amplitude Discrimination and Amplitude Band Width cut more material from the mix, but may also cut out vocals. Higher values make the extraction depend more on the phase of the material and the less on the channel amplitude. Amplitude Discrimination settings between 0.5 and 10 and Amplitude Band Width settings between 1 and 20 work well.

FFT Size Specifies the size of the FFT (Fast Fourier Transform), affecting processing speed and quality. In general, settings between 4096 and 10,240 work best. Higher values (such as the default value of 8192) provide cleaner sounding filters.

Overlays Defines the number of FFTs that overlap. Higher values can produce smoother results or a chorus-like effect, but they take longer to process. Lower values can produce bubbly-sounding background noises. Values of 3 to 9 work well.

Interval Size Sets the time interval (measured in milliseconds) per FFT taken. Values between 10 and 50 milliseconds usually work best, but higher overlay settings may require a different value.

Window Width Specifies the interval (measured as a percentage) used per FFT taken. Values of 30% to 100% work well.

Using the Doppler Shifter effect (Edit View only)

The Doppler Shifter effect creates the increase and decrease in pitch we notice when an object approaches and then passes us, such as when a police car passes with its siren on. The frequency of the noise from the siren starts out at a high pitch and tempo, and it lowers as the car passes you. When the car comes toward you, the sound it makes reaches your ears as a higher frequency because each wave crest is actually compressed by the car moving forward. The first crest leaves the car, and by the time the next one leaves, the car has moved forward, reducing the wavelength of the sound and raising its frequency. The opposite happens as the car passes by: The waves are stretched out, resulting in a lower-pitched sound.

To use the Doppler Shifter effect:

- 1 In Edit View, select an audio range.
- 2 In the Effects tab of the Organizer window, expand Time/Pitch, and double-click Doppler Shifter.
- 3 Set the desired options.

For more information, see [“Doppler Shifter options” on page 144](#).

Doppler Shifter options

The Doppler Shifter effect provides the following options:

Path Type Defines which path the sound source appears to take. Depending on the path type, a different set of options is available.

Note: *Unlike many of the graphs in Adobe Audition effects, the Doppler Shifter graph is noninteractive: You can't directly manipulate the graph. Instead, the graph changes as you adjust the effect's parameters.*

Straight Line lets you specify these options:

- Starting Distance Away sets the virtual starting point (in meters) of the effect.
- Velocity defines the virtual speed (in meters per second) at which the effect moves.
- Coming From sets the virtual direction (in degrees) from where the effect appears to come.
- Passes In Front By specifies how far (in meters) the effect seems to pass in front of the listener.
- Passes On Right By specifies how far (in meters) the effect seems to pass in front of the listener.

Circular lets you specify these options:

- Radius sets the circular dimensions (in meters) of the effect.
- Velocity defines the virtual speed (in meters per second) at which the effect moves.
- Starting Angle sets the beginning virtual angle (in degrees) of the effect.
- Center In Front By specifies how far (in meters) the sound source is from the front of the listener.
- Center On Right By specifies how far (in meters) the sound source is from the right of the listener.

Adjust Volume Based On Distance Adjusts the effect's volume automatically based on the distance values specified.

Adjust Volume Based On Direction Adjusts the effect's volume automatically based on the direction values specified.

Quality Level Provides six different levels of processing quality. Lower quality levels require less processing time, but higher quality levels generally produce better sounding results.

Note: *For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).*

Using chorus, flanger, and phaser effects

These effects can thicken sound or make it outrageous. They range from the Chorus effect's ability to make a single instrument or vocalist sound like a group playing or singing in unison, to the wilder sounds of the Flanger effect and the phaser effects. Although you can apply them in stereo for the most dimensional results, you can use them with mono sound as well.

Using the Chorus effect

The Chorus effect adds richness as if several voices or instruments are played at once. It's a great way to add a degree of presence to a track. You can use it to give a stereo effect to a mono sample (where the left and right channels are identical) or to add harmony or "thickness" to a vocal track. You can also use it to create some truly out-of-this-world special effects.

Adobe Audition uses a direct-simulation method of achieving a chorus effect, meaning that each voice (or layer) is made to sound distinct from the original by slightly varying the timing, intonation, and vibrato. The Feedback setting lets you add extra detail to the result.



You get better results if you convert mono files to stereo before applying the Chorus effect.

To use the Chorus effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Delay Effects, and double-click Chorus.
- 3 Set the desired options.

For more information, see ["Chorus options" on page 145](#).

Chorus options

The Chorus effect provides the following options:

Chorus Characteristics Represent the characteristics used for each voice (layer) in the chorus. While the properties below apply to each voice, they represent ranges of random values, so each voice is unique in each of these characteristics:

- Thickness determines the number of voices to simulate. The final result can end up with an additional voice if Dry Out is set to greater than 0% (causing the original sample to be mixed with the chorused result).

Note: *The more voices you specify, the richer the resulting sound, but the longer the effect takes to be processed.*

- Max Delay specifies the maximum amount of delay allowed. An important component of chorusing is the introduction of short delays (often in the 15-35 millisecond range) that vary in duration over time. If the setting is very small, all the voices start merging into the original, and an unnatural flanging effect might occur. If the setting is too high, a warbled effect might occur, rather like a tape being eaten by a cassette deck.

- Delay Rate determines the time the delay takes to cycle from its zero-to-maximum delay setting. Because the actual delay used varies over time, the pitch of the sample increases or decreases over time, placing each voice slightly out of tune with the others (giving the effect of a separate voice). For example, a value of 2 Hz means the delay could vary from no delay to the maximum delay and back twice per second (sort of a pitch vibrato at twice per second). Note that this setting is a maximum only; if you set it to 2 Hz, the delay might not reach the maximum before it cycles back. If this setting is too low, the individual voices don't vary much in pitch. If it is set too high, the voices may vary so quickly that a warbled effect might occur.
- Feedback includes a percentage of processed chorused voices back into the mix. Feedback can give a waveform an extra echo or reverb effect. A little feedback (less than 10%) can give an extra richness, depending on the delay and vibrato settings. Higher settings produce more traditional feedback, such as loud ringing or other artifacts, that can get loud enough to clip and destroy the signal. Sometimes this clipping is a desired effect, as in the Flying Saucers setting, which generates warbled sounds of UFOs whizzing around your head.
- Spread gives an added delay to each voice, separating them in time by as much as 200 milliseconds (1/5th of a second). High values cause the separate voices to start at different times—the higher the value, the farther apart the onset of each word may be. In contrast, low values cause all voices to be in unison. Depending on other settings, low values can also produce flanging effects, which may be undesirable if your goal is a realistic chorus effect.
- Vibrato Depth determines the maximum variation in amplitude that occurs. For example, you can alter the amplitude of a chorused voice so that it is 5 dB louder or quieter than the original. At extremely low settings (less than 1 dB), the vibrato may be unnoticeable unless the Vibrato Rate is set extremely high. At extremely high settings, however, the sound may cut in and out, creating an objectionable warble. Natural vibratos occur around 2 dB to 5 dB. Note that this setting is a maximum only; the vibrato volume might not always go as low as the setting indicates. This limitation is intentional, as it creates a more natural sound.
- Vibrato Rate determines the maximum rate at which vibrato occurs. With very low values, the resulting voice slowly gets louder and quieter, like a singer that cannot keep his or her breath steady. With very high settings, the result can be jittery and unnatural.



Very high settings can produce interesting special effects (as in the Another Dimension preset).

Stereo Chorus Mode Determines where the individual voices are placed in the stereo field, in addition to how the original stereo signal is interpreted. The Stereo Chorus Mode options, listed here, are active only when you work with stereo files:

- Average Left & Right averages the original left and right channels. If the option isn't selected, the channels are kept separate to preserve any stereo image when you process a stereo source file. (For example, spatial binaural cues such as those that exist in reverberated audio or live stereo recordings are preserved.) Leave this option unselected if the sample was originally monophonic—it won't have any effect except to increase the processing time.

- Add Binaural Cues adds separate delays to the left and right outputs of each voice. This delay can make each voice seem to come from a different direction, but only if there is total separation between what your left ear and right ear hears, as when you listen through headphones. Leave this option unselected for audio that is meant to be played through speakers. In addition, when you add binaural cues, the volume of the right channel for a voice panned all the way to the left is still significant. If no cues are added, no output is sent to the right channel, so greater separation is heard when listening to audio through speakers.
- Narrow Field/Wide Field slider specifies the stereo field, which denotes where in space instruments or other sources are placed within the left and right images of a stereo waveform. The narrower you set the stereo field, the more likely the chorused voices will be near the center of the stereo image. At a setting of 50%, all the voices are spaced evenly along a half circle from left to right. At higher settings, the voices move to the outer edges. If you use an odd number of voices, one is always directly in the center.

Output Lets you specify a mix between the original input (dry) signal and the chorused (wet) signal. Ordinarily, both settings should be less than 100%; otherwise, the overlaying of several voices may cause clipping.

- Dry Out determines how much of the unprocessed signal is mixed into the final output. If you set it to zero, Adobe Audition adds the original voice to the number of processed voices, which are determined by the Thickness setting.



Use the multitrack mixing capabilities of Adobe Audition to dynamically bring in and fade out the chorus. Add the Chorus effect to a copy of the original audio and set Dry Out to zero to create a chorus-only version of the original. In the mixer, insert both the original and the chorus-only version. Use the volume envelope control to adjust the volume of the chorus over time, or tweak the final amplitude of the background chorus with the track's volume settings. This technique is handy for emphasizing portions of singing with a backup chorus.

- Wet Out determines how much of the processed signal is mixed into the final output. In general, lower this value whenever more voices are used. For example, if Thickness is set to 3, a Wet Out value of 40% is appropriate, but if Thickness is set to 10, a Wet Out value of 20% might be better. The best value depends on the number of voices and the desired stereo image field settings.



Keep the original signal (the Dry Out) near 100% and reduce the Wet Out to 30% or so to give the singer or instrument a “backup chorus.”

Highest Quality (But Slow) Ensures the best quality results. Increasing the quality, however, increases the processing time for previewing and applying the effect.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Flanger effect

Flanging was originally achieved by sending an audio signal to two reel-to-reel tape recorders and then physically slowing down the reels of one machine. The resulting sound has a phase-shifted, time-delay effect, characteristic of psychedelic recordings of the 1960s and 1970s. The Flanger dialog box lets you create a similar result by slightly delaying and phasing a signal at predetermined or random intervals.

To use the Flanger effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Delay Effects, and double-click Flanger.
- 3 Set the desired options.

For more information, see [“Flanger options” on page 148](#).

Flanger options

The Flanger effect provides the following options:

Original - Expanded Adjusts the mix between the amount of original (dry) and flanged (wet) signal. You need some of both signals to achieve the characteristic cancellation and reinforcement that occurs during flanging. With Original at 100%, no flanging occurs at all. With Delayed at 100%, the result is a wavering sound, like a bad tape player.

Initial Mix Delay Sets the point in milliseconds at which flanging starts behind the original signal. The flanging effect occurs by cycling over time from an initial delay setting to a second (or final) delay setting.

Final Mix Delay Sets the point in milliseconds at which flanging ends behind the original signal.

Stereo Phasing Sets the left and right delays at separate values, measured in degrees. For example, 180 degrees sets the initial delay of the right channel to occur at the same time as the final delay of the left channel. You can set this option to reverse the initial/final delay settings for the left and right channels, creating a circular, psychedelic effect.

Feedback Determines the percentage of the flanged signal that is fed back into the flanger. With no feedback, the effect uses only the original signal. With feedback added, the effect uses a percentage of the affected signal from before the current point of playback.

Mode Provides three ways of flanging:

- Inverted inverts the delayed signal, causing the waves to cancel out periodically instead of reinforcing the signal. If the Original - Expanded mix settings are set at 50/50, the waves cancel out to silence whenever the delay is at zero.
- Special EFX mixes the normal and inverted flanging effects. The delayed signal is added to the effect while the leading signal is subtracted.
- Sinusoidal makes the transition from initial delay to final delay and back follow a sine curve. Otherwise, the transition is linear, and the delays from the initial setting to the final setting are at a constant rate. If Sinusoidal is selected, the signal is at the initial and final delays more often than it is between delays.

Rate Provides settings for Frequency (measured in Hz), Period (measured in seconds), and Total Cycles (measured in cycles). Each refers to the rate at which the delay cycles between the initial delay and the final delay. Different settings can result in widely varying effects. For example, setting Total Cycles to 0.5 causes the effect to start with the initial delay and end with the final delay. Setting Frequency to 4 causes the flanging cycle to occur four times per second.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Sweeping Phaser effect

Similar to flanging, phasing introduces a variable phase-shift to a split signal and recombines it, creating psychedelic effects first popularized by guitarists of the 1960s. The Sweeping Phaser effect sweeps a notch- or boost-type filter back and forth about a center frequency.

A phase is similar to a flange except that instead of using a simple delay, frequencies are phase-shifted over time. If a phase is used on stereo files, the stereo image can be dramatically altered to create some remarkably interesting sounds.

To use the Sweeping Phaser effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Delay Effects, and double-click Sweeping Phaser.
- 3 Set the desired options.

For more information, see [“Sweeping Phaser options” on page 149](#).

Sweeping Phaser options

The Sweeping Phaser effect provides the following options:

Sweep Gain Lets you vary the amount of gain applied to the phased signal. Take care to avoid clipping when applying higher positive values.

Center Frequency Sets the frequency around which the phase sweeps. Frequencies closer to the middle of the dynamic range of the selected audio produce more dramatic results.

Depth Determines the degree of phasing as a Q value, which is a ratio of width to center frequency. Greater Depth settings make the sweep extend further from the center frequency in both directions (covering a greater frequency range), producing a wider tremolo effect.

Resonance Determines the amount of phase-shifting applied to the signal. You can think of it as a “strength” setting for the phase.

Sweeping Rate Specifies the speed at which the filter sweeps around the center frequency, covering the dynamic range specified by the Depth setting. Values are Hz (cycles per second), Period (milliseconds per beat), and Tempo (beats per minute). To make the sweep occur in time with a song, enter the tempo of the music or a fraction of it. For example, enter 240 for a song with a tempo of 120 to create sweeps in eighth notes.

Stereo Phase Difference Sets the degree at which the sweep interval shifts between the channels of a stereo waveform. You can enter values from –359 to 359. Values farther from 0 or 360 cause the sweep to occur at increasingly distant intervals between the left and right channels. A value of 180 yields a complete difference. Negative numbers are equivalent to their converse positive numbers. For example, –5 and 355 are equivalent, 180 and –180 are equivalent, and 90 and –270 are equivalent.

Sweep Modes Determine the shape of the filter sweep. Sinusoidal and Triangular determine whether the sweep follows a sine wave or a triangle wave. Triangle waves tend to be sharper. Log Frequency Sweep and Linear Frequency Sweep determine whether the sweep occurs in a logarithmic, constant fashion or an even, linear fashion.

Filter Type Determines the type of filter used. Band Pass performs the phase effect around the specified center frequency. Low Pass performs the phase effect from the specified center frequency down. Band Pass is more common since it produces a traditional sound of a sweeping phaser.

Master Gain Adjusts the overall volume output. You can compensate for excessive loss or excessive gain introduced by the effect by entering an overall master gain (measured in decibels) to apply to the resulting audio. A setting of 0 dB is the same as no adjustment.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Graphic Phase Shifter effect

The Graphic Phase Shifter lets you adjust the phase of a waveform by adding control points to a graph.

To use the Graphic Phase Shifter effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Filters, and double-click Graphic Phase Shifter.
- 3 Set the desired options.

For more information, see [“Graphic Phase Shifter options” on page 150](#).

Graphic Phase Shifter options

The Graphic Phase Shifter effect provides the following options:

Phase shift graph Adjusts the phase of a waveform. The horizontal ruler (x-axis) measures frequency, while the vertical ruler (y-axis) displays the degree of phase to shift, where zero is no phase shift.

You can create simulated stereo by creating a zigzag pattern that gets more extreme at the high end on one channel. Put two channels together that have been processed in this manner (using a different zigzag pattern for each) and the stereo simulation is more dramatic. (The effect is the same as making a single channel twice as “zigzaggy”)

Note: For more information about graph controls (such as how to add and remove control points), see the *“Looking at the Work Area”* chapter.



If you adjust the phase of one channel in just a small frequency band, and then use Flange or another effect that depends on phase, you get different results than with audio that isn't phase-shifted.

+/-360° Range Sets the range of the vertical ruler (y-axis) from +360 degrees (top) to -360 degrees (bottom). If this option is unselected, the range is from +180 degrees to -180 degrees.

Log Frequency Scale Sets the values of the horizontal ruler (x-axis) to a logarithmic scale. Use this option to work at finer detail in the lower frequencies. Logarithmic scale is how the ear hears, so each octave occupies a fixed width. If this option is unselected, the x-axis displays values linearly from 0 Hz to the Nyquist (that is, just over twice the sampling frequency), and octaves on the left aren't as wide as octaves on the right.

FFT Size Specifies the size of the FFT (Fast Fourier Transform). Higher sizes usually create more precise results, but they take longer to process.

Channel Specifies the channel(s) to apply the phase shift to.

Note: Process a single channel for the best results. If you apply any phase shift to both channels simultaneously, the resulting file sounds exactly the same when you listen to it in stereo as it did before it was processed.

Flat Resets the graph to its default state.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Changing pitch

The effects in Adobe Audition let you change the pitch, raising or lowering a person’s voice or musical notes. For example, the Pitch Correction effect can correct an out-of-tune vocalist or instrument, and the Stretch effect can stretch or shrink audio without altering pitch or tempo.

Using the Pitch Bender effect (Edit View only)

This effect varies the pitch of the source audio over time. Use the graph to “draw” a tempo to create smooth tempo changes or other effects, such as that of a record or a tape speeding up or slowing down.

To use the Pitch Bender effect:

- 1 In Edit View, select an audio range.
- 2 In the Effects tab of the Organizer window, expand Time/Pitch, and double-click Pitch Bender.
- 3 Set the desired options.

For more information, see [“Pitch Bender options” on page 151](#).

Pitch Bender options

The Pitch Bender effect provides the following options:

Pitch graph Lets you draw a tempo change over time. The horizontal ruler (x-axis) represents duration, while the vertical ruler (y-axis) represents pitch, measured in semitones or beats per minute, depending on the Range option. Points above 0 speed up the sound, while points below 0 slow it down.

Flat Resets the graph to its default state.

Zero Ends Sets the endpoints to no pitch shift. This option is useful if you want to pitch-bend the middle of a selection, so that the endpoints are at the same rate as the surrounding audio.

Quality Level Controls the quality level. Higher quality levels produce the best sound, but they take longer to process. Lower quality levels produce more unwanted harmonic distortion, but they take less time to process. Usually, you won’t notice harmonic distortion for levels from Very Good and higher. Aliasing still occurs, however, when you shift the pitch up, but the higher quality levels greatly reduce the distortion when you shift the pitch down.

Range Sets the scale of the vertical ruler (y-axis) as semitones (there are 12 semitones to an octave) or as beats per minute. For a range in semitones, the pitch changes logarithmically, and you can specify the number of semitones to shift up or down. For a range in beats per minute, the pitch changes linearly, and you must specify both a range and a base tempo. You can specify the exact tempo of a selection to change to different rates, but this isn't required.

Length report Lists what the new length of the file will be. Keep in mind that when decreasing the pitch, very long files can occur, depending on the Range value.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Pitch Correction effect (Edit View only)

The Pitch Correction effect provides two ways to make pitch adjustments for vocals or solo instrumentation. Automatic mode analyzes the audio content and automatically corrects the pitch based on the key you define, without your having to analyze each note. Manual mode creates a pitch profile that you can adjust note-by-note. You can even over-correct vocals to create robotic-sounding effects.

The Pitch Correction effect detects the pitch of the source audio and measures the periodic cycle of the waveform to determine its pitch. The effect can be used on audio that contains a periodic signal (that is, audio with one note at a time, such as for a saxophone, violin, or vocals). Nonperiodic audio, or periodic audio with a high noise floor, can disrupt the effect's ability to detect the incoming pitch, resulting in incomplete pitch correction.

To use the Pitch Correction effect:

- 1 In Edit View, select an audio range.
- 2 In the Effects tab of the Organizer window, expand Time/Pitch, and double-click Pitch Correction.
- 3 Click the Automatic or Manual tab, and set the desired options.

For more information, see [“Pitch Correction options \(Automatic mode\)” on page 152](#) or [“Pitch Correction options \(Manual mode\)” on page 153](#).

Pitch Correction options (Automatic mode)

The Automatic tab of the Pitch Correction dialog box provides the following options:

Reference Channel Specifies which channel to use for source audio. The Pitch Correction effect analyzes and profiles just the periodic signal of the channel you choose, but it applies the pitch correction equally to both channels.

Calibration Specifies the standard pitch calibration for the source audio. In western music, standard pitch calibration is A4 = 440 Hz. Source audio, however, may have been recorded where the calibration was slightly different. In this case, you can raise or lower the Hz value from 430 to 450.

FFT Size Sets the FFT (Fast Fourier Transform) size, or the size of the pieces of data that the effect processes. In general, use smaller values for correcting higher frequencies. For voices, a setting of 2048 or 4096 sounds most natural, and a setting of 1024 creates robotic effects.

Scale Specifies the scale type that best suits the material. You can choose from Major, Minor, or Chromatic. Chromatic works best if only a slight correction is needed to “pull” the pitch of any tone to the closest chromatic note. Major and Minor allow for larger corrections, such as where the source audio is more than a single half-step interval off pitch from the desired musical scale.

Key Sets the key for the corrected material. This option is available only if Scale is set to Major or Minor (since the Chromatic scale includes all 12 tones and isn’t key-specific). Typically, the key is the same as the one intended for the source audio.

Attack Governs how quickly Adobe Audition corrects the pitch toward the scale tone. Faster settings are usually best for audio made up of short notes, such as a fast passage played by a trumpet. An extremely fast attack can also achieve a robotic quality. Slower settings result in more natural-sounding correction on longer sustaining notes, such as a vocal line where the singer holds notes and adds vibrato. Because source material can change throughout a musical performance, you can get the most natural-sounding results by correcting short pieces (such as individual phrases) at a time.

Sensitivity Defines a threshold beyond which a note isn’t corrected. Sensitivity is measured in cents, and there are 100 cents per semitone. For example, a Sensitivity value of 50 cents means a note must be within 50 cents (half a semitone) of the target scale tone before it is corrected automatically.

Correction Meter Displays the correction as it occurs during preview, showing you when flat tones are raised and sharp tones are lowered, and by how much.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Pitch Correction options (Manual mode)

The Manual tab of the Pitch Correction dialog box provides the following options:

Reference Channel Specifies which channel to use for source audio. The Pitch Correction effect analyzes and profiles just the periodic signal of the channel you choose, but it applies the pitch correction equally to both channels.

Calibration Specifies the standard pitch calibration for the source audio. In western music, standard pitch calibration is A3 = 440 Hz. Source audio, however, may have been recorded where the calibration was slightly different. In this case, you can raise or lower the Hz value from 430 to 450.

FFT Size Sets the FFT (Fast Fourier Transform) size, or the size of the pieces of data that the effect processes. In general, use smaller values for correcting higher frequencies. For voices, a setting of 2048 or 4096 sounds most natural, and a setting of 1024 creates robotic effects.

Pitch Reference graph Displays the pitch in single half-step increments along the vertical ruler (y-axis) and time along the horizontal ruler (x-axis), so that you can see the exact transitions of the pitch in all parts of the source audio.

Pitch Profile and Pitch Correction lines Show the amount of correction you need at any given point in a waveform. The red Pitch Profile line is not editable, and the green Pitch Correction line appears as you modify the edit envelope in the Pitch Edit graph.

Pitch Edit graph Displays the pitch in half-step pitch increments along the y-axis and time (in the same format as the timeline along the bottom of the display window) along the x-axis. If the source is a stereo file, the graph shows the channel you specified as the Reference Channel.

The Pitch Edit graph display area includes several components:

- The adjustable blue envelope line lets you control how pitch correction is applied. Click on the line to create a control point, which you can drag to set the correction for the audio material at that point in the time scale. For more precision, right-click a control point to display the Edit Point dialog box, where you can enter values for Time Index (x-axis) and Pitch (y-axis).
- A readout below the Pitch Edit graph shows the position of a control point as you drag it.
- A playback cursor moves through the waveform when you listen to it.

Vertical portion bar Lets you scroll to a different portion of the Pitch Reference graph. Zoom in or out by dragging the edge of the slider.

Horizontal portion bar Lets you scroll to a different time section of the waveform in both graphs. Zoom in or out by dragging the edge of the slider.

Note: To use the portion bars, you can also zoom out by right-clicking the slider and choosing *Zoom Out* or *Zoom Out Full*, and you can right-click-zoom for more precision.

Zoom buttons Zoom into and out of the Pitch Edit graph. The zoom buttons always keep the “0” line at the vertical center. The maximum amount you can zoom in is plus or minus one-half step. The maximum you can zoom out is plus or minus 200 cents. The waveform always retains its same vertical size and isn’t affected by the vertical zoom.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Stretch effect

The Stretch effect lets you change the pitch of an audio signal, the tempo, or both. For example, you can use the effect to transpose a song to a higher key without changing the tempo, or you can use it to slow down a passage without changing the pitch. You can also vary pitch and tempo over the length of the audio, giving the effect of raising and lowering pitch or slowing down and speeding up the tempo.

To use the Stretch effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Time/Pitch, and double-click Stretch.
- 3 Set the desired options.

For more information, see [“Stretch options” on page 154](#).

Stretch options

The Stretch effect provides the following options:

Constant Stretch, Gliding Stretch Performs a constant stretch, in which there is no change in the amount of stretch on the selection, or a gliding stretch, which varies the amount of stretch from an initial percentage to a final percentage.

Stretch %, Initial %, Final % Sets the amount of stretch applied to the waveform (less than 100% compresses the wave). If you use Gliding Stretch, sliders for both initial and final settings are available. Changes in the sliders are reflected in each of the Ratio and Length boxes. Specify separate values for Initial and Final to stretch the waveform in a linear fashion from one ratio to another.

Ratio, Length Specifies the ratio (in percentage) and final length (in time) for the stretch. Specifying a value for one automatically changes the other. If the initial and final lengths are different, then the actual final length will be exactly (initial+final)/2 when in Preserve Pitch mode.

Transpose Lists the musical transposition amounts. The corresponding numerical values are entered into the stretch sliders automatically. For example, to transpose sound up one semi-tone (one half-step on a keyboard) choose 1# for one sharp.

Precision Defines the overall faithfulness to the sound's quality, with higher quality taking longer to process. 8-bit or low-quality audio files can be processed quickly with the Low Precision setting, whereas a professionally-recorded audio file may require stretching using the High Precision setting.



A quick way to determine which precision quality to use is to process a small portion of the audio at each setting until you find the best balance of quality and processing time.

Stretching Mode Provides three stretching options:

- Time Stretch (Preserves Pitch) lets you decrease and increase the tempo without changing the pitch. Lower percentages slow the tempo and higher percentages speed it.

Note: Use this setting to make a 33- or 28-second commercial exactly 30 seconds.

- Pitch Shift (Preserves Tempo) lets you raise and lower the pitch without changing the tempo. Lower percentages raise the pitch and higher percentages lower it.



Use this setting to make a voice sound deeper or higher without affecting the original playback speed. Or, use differing initial and final percentages to raise and lower the pitch without affecting the tempo.

- Resample (Preserves Neither) lets you change both the pitch and tempo. Percentages below 100 increase the tempo and raises the pitch, while percentages above 100 decrease the tempo and lower the pitch.

Pitch And Time Settings Provide the following options:

- Splicing Frequency determines the size of the chunk of audio data used when you preserve pitch or tempo while elongating or truncating a waveform. The higher the value, the more precise the placement of stretched audio over time. However, artifacts are more noticeable as rates go up. At higher precision, lower splicing frequencies may add stutter or echo. If the frequency is too high, sound might be tinny or voices might have a tunnel-like quality.



If Low Precision mode is used, you can improve the quality of stretched monotonal (pure tone) samples by choosing a splicing frequency that's evenly divisible into the frequency of the sample. Use the Frequency Analysis window to find the sample's base frequency, and then divide by an integer to get the splicing frequency. For example, if the tone is 438 Hz, dividing by 20 gives 21.9 Hz. Thus, a splicing frequency of 21.9 Hz will greatly improve quality by reducing phase artifacts. For nontonal or noisy samples, the splicing frequency doesn't matter as much.

- Overlapping determines how much the current chunk of audio data overlaps with the previous and next ones. (When stretching or compressing audio, chunks are overlapped with previously transformed chunks.) If stretching produces a chorus effect, lower the Overlapping percentage. Doing so, however, can produce a choppy sound. Simply adjust Overlapping to strike a balance between choppiness and chorusing. Overlapping can be as high as 400%, but you should use this value only for very high speed increases (200% or more).
- Choose Appropriate Defaults applies good default values for Splicing Frequency and Overlapping. This option is good for preserving pitch or tempo.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Adding delays and echoes

Delay refers to separating copies of an original signal by some number of milliseconds. *Echoes* are sounds that are delayed far enough in time so that you hear each as a distinct copy of the original sound. Both delays and echoes are a great way to add ambiance to a track where reverb or chorusing might muddy the mix.

Using the Delay effect

Delay can be used to create single echoes, as well as a number of other effects. Delays of 35 milliseconds or more create discrete echoes, while those between 15-34 milliseconds can create a simple chorus or flanging effect. (These results won't be as effective as the actual Chorus or flanging effects in Adobe Audition, as the delay settings are fixed and don't change over time.)

By further reducing a delay to between 1 and 14 milliseconds, you can spatially locate a mono sound (which has the same information in both the left and right channels) so that the sound seems to be coming from the left or the right side, even though the actual volume levels for left and right are identical.

To use the Delay effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Delay Effects, and double-click Delay.
- 3 Choose Effects > Delay Effects > Delay.
- 4 Set the desired options.

For more information, see [“Delay options” on page 156](#).

Delay options

The Delay effect provides the following options:

Delay Adjusts the delay for both the left and right channels from -500 milliseconds to +500 milliseconds. Entering a negative number means that you can move a channel ahead in time instead of delaying it. For instance, if you enter 200 milliseconds for the left channel, the delayed portion of the affected waveform is heard before the original part.

Mixing Sets the percentage of Delayed (wet) signal and unprocessed Original (dry) signal to be mixed into the final output. A value of 50 mixes the two evenly.

Invert Changes the positive values of the selected waveform to negative values. Inverting the delayed signal can be used for special effects, such as creating a quick-and-dirty comb filter. Cancellation occurs when you mix an inverted waveform with the original.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Dynamic Delay effect

The Dynamic Delay effect lets you change the amount of delay over the length of a waveform. For example, you could set a 2 millisecond delay for the first five seconds of audio, a 20 millisecond delay for the next 15 seconds, a 7 millisecond delay for the next 10 seconds, and so on.



Dynamic Delay is especially cool when used as a real-time effect in Multitrack View. If you add the dynamic delay (or Dynamic EQ, which has a similar principle) to Multitrack View, you get a new envelope that determines the delay.

To use the Dynamic Delay effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Delay Effects, and double-click Dynamic Delay.
- 3 Set the desired options.

For more information, see [“Dynamic Delay options” on page 157](#).

Dynamic Delay options

The Dynamic Delay effect provides the following options:

Original and Delayed Mixes in the desired amounts of original and delayed audio.

Invert Inverts the delayed signal, causing the waves to cancel out periodically, instead of reinforcing the signal.

Delay graph Indicates the amount of delay. The horizontal ruler (x-axis) represents time, and the vertical ruler (y-axis) represents the length of delay.

Feedback graph Indicates the amount of feedback. The x-axis represents time, and the y-axis represents the percentage of feedback.

Loop Graphs Makes the graphs apply only as long as specified, so they loop repeatedly over the course of the selection. With this option selected, you can specify how fast the cycle should move (Frequency), how long the graphs should be (Period), and how many times you want the graph to loop in the given selection (Total Cycles). Changing one of these values changes the other two. If Loop Graphs is unselected, the graphs are the same length as the selection, so the dynamic changes to the delay are graphed across the whole selection.

Stereo Curve Delay (Edit View), Stereo Curve Difference (Multitrack View) Represents the number of milliseconds one channel is behind with respect to the other when following the envelope you draw in Multitrack View. Positive values delay the right channel, while negative values delay the left. The audio itself isn't delayed, just the curve used. As a result, the phasing sound of the delay lags in one channel with respect to the other.

Automated (Multitrack View) Applies values automatically. Select this option in the Delay section or Feedback section of the dialog box to draw the Delay curve or Feedback curve using a purple envelope on waveform clips. Note that you must have View > Show FX Parameter Envelopes In Multitrack View enabled to use the Automated option.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Echo effect

This effect adds a series of repeated, decaying echoes to a sound. (For a single echo, use the Delay effect instead.) You can create effects ranging from a Grand Canyon-type “Hello-ello-llo-lo-o” to metallic, clanging drainpipe sounds by varying the delay amount. By equalizing the delays, you can change a room’s characteristic sound from one with reflective surfaces (creating echoes that have bright, shiny, high-end sounds) to one that is almost totally absorptive (meaning very few high-end sounds are reflected).

Note: Make sure that enough silence is at the end of the waveform for the echo to end. If the echo is cut off abruptly before it fully decays, undo the Echo effect, add several seconds of silence by choosing Generate > Silence, and then reapply the Echo effect.



You can create striking stereo echo effects by setting different left and right values for the Decay, Delay, and Initial Echo Volume controls.

To use the Echo effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Delay Effects, and double-click Echo.
- 3 Set the desired options.

For more information, see [“Echo options” on page 158](#).

Echo options

The Echo effect provides the following options:

Echo Characteristics Lets you define these echo characteristics.

- Decay determines the falloff ratio of an echo. Each successive echo tails off at a certain percentage less than the previous one. A decay setting of 0% results in no echo at all, while a decay of 100% produces an echo that never gets quieter.
- Delay specifies the number of milliseconds between each echo. For example, a setting of 100 milliseconds results in a 1/10th-second delay between successive echoes.
- Initial Echo Volume sets the percentage of echoed (wet) signal to mix with the original (dry) signal in the final output.
- Lock Left/Right links the sliders for Decay, Delay, and Initial Echo Volume, maintaining the same settings for each channel.
- Echo Bounce makes the echoes bounce back and forth between the left and right channels. If you want to create one echo that bounces back and forth, select an initial echo volume of 100% for one channel and 0% for the other. Otherwise, the settings for each channel will bounce to the other, creating two sets of echoes on each channel.

Successive Echo Equalization Provides an eight-band echo “quick filter” for specifying which frequencies are removed from an echo first. Each successive echo is passed back through the equalizer, letting you simulate the natural sound absorption of a room. A setting of 0 leaves the frequency band unchanged, while a maximum setting of -15 decreases that frequency by 15 dB. And, because -15 dB is the difference of each successive echo, some frequencies will die out much faster than others.

Note: *Unlike most other equalizer-like controls, you can't increase the attenuation of frequencies with Successive Echo Equalization; you can only decrease them.*

Continue Echo Beyond Selection Continues the effect beyond the right boundary of the waveform's selected range. In other words, the echo decays naturally beyond the boundary, while unselected sounds aren't echoed. Leaving this option unselected makes the echo stop at the right boundary visible in the waveform window, meaning that if the window is zoomed in, the echoing stops before the range or file ends.

Note: *For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).*

Using the Echo Chamber effect

The Echo Chamber effect can simulate the ambiance of almost any room. Settings let you specify a virtual room's size and surface characteristics, along with the placement of virtual microphones. The number of echoes is adjustable up to 500,000. Keep in mind that the more echoes you include, the more time Adobe Audition needs to process the effect.

You can create a spatial, stereo expansion effect by setting the virtual microphones farther apart than your actual stereo speakers. For example, if your stereo speakers are 6 feet apart, try setting the left and right virtual microphones 20 or 30 feet apart.



Make sure that enough silence is at the end of the waveform for the echo to end. If the echo is cut off abruptly before it fully decays, undo the Echo effect, add several seconds of silence by choosing **Generate > Silence**, and then reapply the Echo Chamber effect.

To use the Echo Chamber effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Delay Effects, and double-click Echo Chamber.
- 3 Set the desired options.

For more information, see [“Echo Chamber options” on page 159](#).

Echo Chamber options

The Echo Chamber effect provides the following options:

Room Size (Feet) Specifies the width, length, and height of the virtual room, measured in feet. (A foot is approximately 0.3 meters.) Room sizes can be as large as you like, but memory requirements increase as room size increases.

Intensity Determines the percentage of amplitude of the direct (original) signal. Because echoes (picked up by the virtual microphone) add to a signal's overall amplitude, you should always set Intensity to less than 100% to avoid clipping. In fact, the more echoes there are, the lower the percentage should be. For example, a setting of about 30% is appropriate for 100 echoes, whereas 15% is better for 1,000 echoes.

Echoes Specifies the number of echoes to produce. To achieve a nice, reverberating, ambiance effect, use at least 10,000 echoes. The more echoes that are generated, the truer the result will sound, but the longer the processing will take. Using 25,000 echoes produces a very realistic result.



To speed up testing of a virtual chamber's size and overall room sound, you may want to generate only 100 echoes. Once you achieve the sound you want, you can undo the test effect and increase the number for final production. A very fast system should be able to generate up to 500,000 echoes, depending on the virtual room size and available memory.

Damping Factors Describes the type of virtual room in which the audio is being played by letting you select how much the sound is dampened by each wall, the floor, and the ceiling. These factors can simulate wall coverings, floor coverings, and other objects in the room that absorb sound. A high damping factor is reflective (like cement), while a low damping factor absorbs sound (like carpeting and sound-proofing panels). Although all frequencies are absorbed equally (unlike in real life), the result is more realistic than if you use only the basic echo settings.

Signal And Microphone Placement (Feet) Provides the following options:

- **Source Signal** specifies the distance (Dx) of the source signal (the waveform or selection) from the left wall, the back wall, and above the floor of the virtual room. (For stereo waveforms, you can set different values for the left and right channels.) The signal then simulates a single, nondirectional point source, meaning that the sound radiates outward in all directions. The distance between the source and the walls affects which frequencies are enhanced and is crucial to the overall ambient effect. If you enter a value that exceeds the dimensions of the room, Adobe Audition uses the greatest possible value based on the dimensions.
- **Microphone** specifies the distance (Dx) of the virtual microphone from the left wall, the back wall, and above the floor. (For stereo waveforms, you can set values for two virtual microphones.) The resulting echoes emulate what the microphone would pick up in the room at the specified location. In a stereo setting, place the virtual microphones one foot apart to simulate human ears. The placement of and distance between the microphones gives the brain cues about the directions of each echo and the size of the room. Try listening with headphones to virtual microphones that have been placed far apart; the sound has a very large aural or spacey feel.

Note: Make sure that you always place the virtual microphones sufficiently far away from the source. If a microphone and source are too close, you will hear only the source and no echo. This scenario is analogous to placing your ear right next to the sound source, where you hear only the sound because of its loudness.



Give a stereo effect to monaural audio by adding ambiance: Set the left microphone one or two feet from the right microphone to simulate a listener's ears and give the effect of "being there," especially effective with headphones. Make sure that you first convert the mono signal to stereo (choose Edit > Convert Sample Type) so that you can choose separate virtual microphone locations.

Mix Left/Right Into Single Source Combines the left and right channels of a stereo waveform before processor occurs. Select this option for faster processing, but leave it unselected for a fuller and richer stereo effect.

Damping Frequency Specifies the upper frequency limit of reflected sound. For instance, if you set Damping Frequency to 7000 Hz, frequencies above 7000 Hz are cut for each unit of time. Use lower values for a warmer sound.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Multitap Delay effect

Multitap Delay can be thought of as a combination of the Delay, Echo, Filter, and Reverb effects. You can create up to 10 delay units, each with its own delay, feedback, and filtering settings.

If one delay unit is placed inside another (as viewed in the chart above the controls), then the echo occurs more than once. As audio travels down the delay line (represented in the chart by the bottom horizontal arrow pointing to the right) portions at any point can be fed back into the delay line anywhere behind the given offset and at any feedback amount, with any high or low cut filter. Experiment to achieve some very interesting effects.

Each delay unit is represented in the graph as a back-leading arrow starting at the Offset and going back the number of milliseconds stated under Delay. A single delay unit is much the same as the Echo function, but with a slightly different filtering setup. (It uses two sliding bands with variable cutoff points instead of eight bands of filtering.)

To use the Multitap Delay effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Delay Effects, and double-click Multitap Delay.
- 3 Set the desired options.

For more information, see [“Multitap Delay options” on page 161](#).

Multitap Delay options

The Multitap Delay effect provides the following options:

Delay Units Displays the different delay units and their settings in the format: Delay [delay time] at [Offset] ([Feedback percentage]). Select a delay unit in the list to adjust its delay settings. Click Add New to create a new Delay Unit with the current Delay Settings. Click Remove to delete the selected Delay Unit.

Delay Lets you adjust the following delay parameters for each unit:

- Offset adjusts the point in the delay line from which Adobe Audition takes the audio. It is then mixed into an earlier point in the delay line, which causes echoing. Keep in mind that it's the relative positions of the offsets of the delay units that make a difference, not their absolute position. For instance, if you have two delay units at offsets of 200 and 500, the resulting audio sounds the same if they were at 100 and 400. The difference between the offsets is what is important.
- Delay specifies the number of milliseconds to wait before feeding the audio back into the delay line. The result is an echo with a period of the delay given to be generated. With several delay units of varying delays added, the final echo pattern can become very complex. Very short delays give ringing or robotic sounding events. Longer delays give more distinct echoes.

- Feedback represents the percentage of the original signal to feed back into the delay line. If the feedback is set too high, ringing and feedback occur. The audio gets louder and louder until it clips and becomes distorted. However, sometimes you may want this effect, which is similar to the feedback you hear when a live microphone is set too close to a loudspeaker. If the feedback percentage is too low, then not very much of the original signal is fed back into the loop. This results in a very subtle effect.
- Allpass Feedback helps prevent the DC component from getting out of hand (the waveform tending upwards or downwards until it clips). With the option selected, audio from the destination of the delay loop is mixed back into audio from the originating delay offset. Instead of going one way (from the offset back a certain number of milliseconds), it also goes from the destination up to the source, creating a sort of forward feedback, or “feed-forward.” This setting is handy when designing reverb effects.

Low-Cut Filter and High-Cut Filter Filter the audio being fed back into the delay line. Low-Cut Filter reduces or boosts the low frequencies, depending on the Cutoff and Boost settings. High-Cut Filter reduces or boosts the high frequencies, making each successive echo filtered slightly differently for interesting effects. The Low-Cut Filter and High-Cut Filter provide two options:

- Cutoff determines which frequencies are affected. Frequencies below this setting are affected by the Low-Cut Filter. Frequencies above this setting are affected by the High-Cut Filter. Changes in the cutoff value affect the tone of the echoes, as more or less of the frequencies are affected by the filter.
- Boost sets the amount of filtering. Boost settings are usually negative, which means the audio is reduced in the affected frequency range. Lower negative values result in more audio being cut. Positive values result in boosted frequencies.



When designing a reverb, cut some of the high frequencies to simulate their absorption by surrounding walls. In addition, when echoing, frequencies generally aren't boosted. However, you can create interesting effects by entering positive values. Boosting a Low-Cut Filter while reducing the feedback setting is identical to reducing a High-Cut Filter and increasing the feedback setting.

Channel buttons Provide the following options for use with stereo source audio. To spread out the effect in the stereo domain, at least one delay unit in the group should be cross-channel or single-channel only:

- Left Only makes the delay appear only on the left channel of stereo audio.
- Right Only makes the delay appear only on the right channel of stereo audio.
- Discrete Stereo makes the delay appear on both channels of stereo audio.
- Stereo Swap maps the delay for the left channel to the right channel and vice-versa.
- From Left to Right delays the audio from the left channel and puts it onto the right channel.
- From Right to Left delays the audio from the right channel put it onto the left channel.
- From Center to Surround delays the audio from the center channel (the audio common to both left and right channels) and puts it onto the surround channel (the left channel inversed with respect to the right channel).

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Adding reverb

When a sound occurs, it bounces off of different surfaces on its way to your ears. For example, when someone sings in a room, that sound is reflected off the walls, ceiling, and floor, as well as any objects in the room. This reflected sound is called *reverberation*, or *reverb* for short. All these reflected sounds might reach your ears so closely together that you cannot discern them as separate echoes. However, they give an impression of space. With Adobe Audition, you can customize the reverb and replicate a variety of room environments.



For the most precise control of an effects mix in Multitrack View, set real-time reverbs to 0% Original and 100% Reverb. Then, use the effects mixer to control the ratio of dry to reverberant sound.

Using the Full Reverb effect

Full Reverb, like the standard Reverb effect, simulates acoustic space. It's also convolution-impulse-based (like standard Reverb), meaning no ringing, metallic, or other artificial sounding artifacts are present. However, specific resonance can be achieved if desired. (See [“Full Reverb options \(Coloration tab\)” on page 165.](#))

The Full Reverb effect has some unique features, such as Perception, which simulates room irregularities, and source location to place the “singer” off-center, and have the early reflections realistically model their position within rooms that have acoustically desirable dimensions that you can customize. Practically any wall surface or other sound-affecting factors can be simulated by changing the reverb's frequency absorption by using a three-band, parametric-EQ style interface (in the Coloration tab).

Note: *Because the Full Reverb effect can take longer to process than the other effects, it may not be the best choice for using in real time in Multitrack View. If you use the Full Reverb effect on a track, consider locking the track afterwards so that it doesn't slow down your editing process.*

To use the Full Reverb effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Delay Effects, and double-click Full Reverb.
- 3 Click the General Reverb tab, the Early Reflections tab, or the Coloration tab, and specify the options you want on each.

Note that when any of the reverb characteristics are modified, a new impulse is built to simulate the environment you specify. (An “impulse” is the data by which every other sample in a waveform is multiplied.) The impulse can be several megabytes in size, requiring more CPU processing power, so you might have to wait a few seconds after clicking Preview for the reverb to be built. The results, however, are much more natural sounding and easier to tailor. Once built, the preview generally runs in real time, and subsequent previews don't require rebuilding the impulse, nor does adjusting any of the Mixing options or selecting Include Direct.

- 4 Specify any Mixing options you want.

For more information, see [“Full Reverb options \(General Reverb tab\)” on page 164](#), [“Full Reverb options \(Early Reflections tab\)” on page 164](#), and [“Full Reverb options \(Coloration tab\)” on page 165](#).

Full Reverb options (General Reverb tab)

The General Reverb tab provides the following basic options for controlling reverb:

Total Length Specifies the overall number of milliseconds the reverb takes to decay 60 dB. However, depending on the Coloration parameters, certain frequencies may take longer to decay to 60 dB, while other frequencies may decay much faster. Longer values give longer reverb tails, but they also take longer to be computed. The effective limit is about 6000 milliseconds (a 6-second tail). The actual tail generated is much longer than this in order to allow for decaying into the background noise level.

Attack Time Specifies the number of milliseconds the reverb takes to build to its maximum amplitude. Generally, reverbs tend to build up over a short time span, and then decay at a much slower rate. Interesting effects can be heard with extremely long attack times (like 400 milliseconds or more).

Diffusion Controls the rate of echo buildup and how diffuse the echoes are. High diffusion values (above 900) give very smooth reverbs, without distinct “echoes” heard in them. Lower values produce more distinct echoes since the initial echo density is lighter, but the density builds over the life of the reverb tail.



Interesting “bouncy” effects can be obtained by using low Diffusion values and high Perception values. Using low Diffusion values, and somewhat low Perception values with long reverb tails, gives the effect of a football stadium or similar arena.

Perception Models irregularities in the environment (objects, walls, connecting rooms, and so on). Low values create a smoothly decaying reverb without any frills. Larger values give more distinct echoes (coming from different locations).



If a reverb is too smooth, it may not sound natural. Values up to about 40 give just enough variation to the reverb to simulate small room variations.

Set Reverb Based On Early Reflection Room Size Adjusts the Total Length and Attack Time settings to roughly match the room size specified in the Early Reflections tab and produce a more convincing sounding reverb. You can then fine-tune the length and attack time to change the effect.

Full Reverb options (Early Reflections tab)

The Early Reflections tab provides the following options to control the size and shape of a virtual room:

Room Size Sets the volume of the virtual room, as measured in cubic meters. The larger the room, the longer the reverbs. Use this control to create virtual rooms from only a few meters to giant coliseums.

Dimension Specifies the ratio between the room’s width (left to right) and depth (front to back). A sonically appropriate height is calculated and reported as in the dialog box as Actual Room Dimensions. Generally, rooms with width-to-depth ratios between 0.25 and 4 provide the best sounding reverbs.

Left/Right Location Lets you place the source off-center to produce a different set of incoming early reflection echoes. Selecting the Include Direct option in the Mixing section also adjusts the Original Signal to sound as if it is coming from the same location by delaying one of the channels. Very nice effects are possible with singers being very slightly off center, say, between 5% and 10% left or right.

High Pass Cutoff Prevents the loss of low-frequency (100 Hz or less) sounds, such as bass or drums. These sounds can get phased out when using small rooms if the early reflections mix with the original signal. Specify a frequency above that of the sound you wish to keep. Good settings are generally between 80 Hz and 150 Hz. If the cutoff setting is too high, you may not get a realistic image of the room size.

Set Reverb Based On Early Reflection Room Size Sets an appropriate reverb length and attack time to match the specified the room size, making for a more convincing listening experience. If desired, you can adjust the reverb length and attack later.

Full Reverb options (Coloration tab)

The Coloration tab provides the following options for filtering or “coloring” the quality of the reverb:

Amplitude/dB sliders Provide a quick way to get a different reverb. They are located to the right of the graph and specify the low shelf, mid band, and high shelf bands, from left to right.



If you want to enhance a quality of audio, such as a singer’s voice, try boosting the frequencies just around the natural frequency of the voice to enhance resonance in that range (for example, from 200 Hz to 800 Hz).

Low Shelf, Mid Band, and High Shelf Specifies the cornering frequency for the shelves or the center frequency for the mid band. For example, to increase reverb warmth, lower the high shelf frequency while also reducing its amplitude.

Q Sets how wide the mid band’s affective area is. Lower values affect a narrower range of frequencies and higher values affect a wider range.



For distinct resonance, use values like 10 or higher. For general boosting or cutting a wide range of frequencies, use lower values, like 2 or 3.

ms Specifies the number of milliseconds the reverb takes to decay at each frequency following the coloration curve. Values up to 700 work fine. For more colored reverbs, use lower settings (such as 100 to 250). Basically, the lower the value, the more the graph affects the resultant reverberation.

Original Signal (Dry) Adjusts the amount of original signal with respect to the other levels to create a sense of distance between the listener and the source. Use a low signal level to sound far away. Use a high level (near 100%) along with low levels for the other settings to create a sense of being close to the source.

Early Reflections Controls the percentage of echoes that first reach the ear, giving a sense of the overall room size. Too high a value can result in an artificial sound, while too low a value can lose the audio cues for the room’s size. Half the volume of the original signal is a good starting point.

Reverb (Wet) Controls the volume of the dense layer of sound usually associated with reverb. Like the Early Reflections setting, this option should be fine-tuned to give a pleasing mix. If it’s too loud, the reverb sounds very unnatural. The balance between the reverb and the original signal gives the impression of distance, so increase the reverb volume with respect to the original signal to make the source sound farther away.

Include Direct Phase-shifts (delays) the original signal's left and right channels slightly to match the direction of the incoming early reflections. The Left/Right Location setting in the Early Reflections tab determines the direction the audio appears to come from when listened to with stereo headphones.

Combine Source Left And Right Combines the left and right channels of a stereo waveform before processing occurs. Select this option for faster processing, but leave it unselected for a fuller and richer stereo effect.



Select this option if both channels are identical (that is, if they originated from a monophonic sample).

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the QuickVerb effect

Like Full Reverb and Reverb, the QuickVerb effect adds reverberation to audio to simulate a different acoustic space. It is faster to use, however, because it isn't convolution-based like Full Reverb and Reverb (both of which increase the processing load on your system). As a result, you can make real-time changes more quickly and effectively in Multitrack View, without needing to “lock” effects to a track. For slightly faster processing and more control, you can also use the Studio Reverb effect. For more information, see [“Using the Studio Reverb effect” on page 168](#).

To use the QuickVerb effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Delay Effects, and double-click QuickVerb.
- 3 Set the desired options.

For more information, see [“QuickVerb options” on page 166](#).

QuickVerb options

The QuickVerb effect provides the following options:

Room Size Sets the room size.

Decay Adjusts the amount of reverberation decay in milliseconds.

Diffusion Simulates the absorption of the reverberated signal as it is reflected off of surfaces, such as carpeting and drapes. Lower settings create more echoes, while higher settings produce a smoother reverberation with fewer echoes.

High Cutoff Freq Specifies the highest audio frequency at which reverb can occur.

Low Cutoff Freq Specifies the lowest audio frequency at which reverb can occur.

Original Signal (Dry) Sets the percentage of source audio to leave in the effect.

Reverb (Wet) Sets the percentage of reverb to place in the effect.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Reverb effect

The Reverb effect lets you simulate acoustic space, and it consists of both early reflections and echoes that are so closely spaced that they're perceived as a single decaying sound. The Reverb effect is different from the basic Echo effect in that the delays aren't repeated at regularly spaced intervals.

The Reverb effect can create a wide range of high-quality reverb results. It can reproduce acoustic or ambient environments such as a coat closet, a tiled bathroom shower, a concert hall, or a grand amphitheater. The echoes can be spaced so closely together and made to occur at such random times that a signal's reverberated tail decays smoothly over time, creating a warm and natural sound. Alternatively, initial early-reflection delays can be used to give a sense of room size, depending upon the initial delay times.

The difference between the Reverb effect and the Full Reverb effects is that Full Reverb is newer, and it provides more options and better audio rendering. However, you may prefer the older Reverb effect if that's what you're used to using.

Note: *Because the Reverb effect can take longer to process than the other reverb effects, it may not be the best choice for using in real time in Multitrack View.*



To simulate rooms that have both echoes and reverb, use the Echo effect first to establish the "size" of the room sound, and then use the Reverb effect to make the sound more natural. This technique can create a sense of spaciousness in a monophonic signal (one that has been recorded as or converted into a stereo audio file). Even a Total Reverb Length as little as 300 milliseconds can open up the perceived spaciousness of a dry sound (one that was recorded without any effects or reverb).

To use the Reverb effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Delay Effects, and double-click Reverb.
- 3 Set the desired options.

For more information, see ["Reverb options" on page 167](#).

Reverb options

The Reverb effect provides the following options:

Total Reverb Length Sets the number of milliseconds the reverb takes to tail off to infinity (about -96 dB). Use values below 400 for small room sizes, use values between 400 and 800 for medium room sizes, and use values above 800 for very large room sizes, such as concert halls. For example, enter 3000 milliseconds to create reverb tails for a giant amphitheater.

Attack Time Sets the amount of time it takes the reverb to gain full strength. For short reverb times, the attack time should be smaller. In general, a value about 10% as long as the Total Reverb Length works well. However, you can create interesting and subtle effects by using longer attack times with shorter reverb lengths, and conversely, by combining very short attack times with long reverb lengths.

High Frequency Absorption Time Simulates natural absorption so that high frequencies are reduced (attenuated) as the reverb decays. Faster absorption times simulate rooms that are occupied and have furniture and carpeting, such as night clubs and theaters. Slower times (especially over 1000 milliseconds) simulate rooms that are emptier, such as auditoriums, where higher frequency reflections are more prevalent. In acoustic environments, higher frequencies tend to be absorbed faster than lower frequencies.

Perception Adds subtle qualities to the environment by changing the characteristics of the reflections that occur within a room. Lower values create smoother reverb without as many distinct echoes. Higher values simulate larger rooms, cause more variation in the reverb amplitudes, and add spaciousness by creating distinct reflections over time.



A setting of 100 and a reverb length of 2000 milliseconds or more creates interesting canyon effects.

Original Signal (Dry) Sets the percentage of source audio to leave in the effect. In general, the more reverb you add, the lower the volume of the original signal. In most cases, 90% works well.



To add spaciousness to an instrument, keep the dry signal higher, or at 100%. If you're trying to achieve a special effect with reverb, you might want to reduce the volume of the original signal. If the reverb is so great that audio begins to clip, try reducing both the dry and the reverberated signal strength.

Reverb (Wet) Sets the percentage of reverb to place in the effect. To add spaciousness to a track, keep the percentage of reverb lower than the percentage of the original signal. However, you may want to increase the percentage to simulate physical distance from the audio source (where reverb is heard in greater proportions to the original signal).

Combine Source Left And Right Combines the left and right channels of a stereo waveform before processing occurs. Select this option for faster processing, but leave it unselected for a fuller and richer stereo effect.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Studio Reverb effect

Like Full Reverb, QuickVerb, and Reverb, the Studio Reverb effect adds reverberation to audio to simulate a different acoustic space. It is faster to use than Full Reverb and Reverb, however, because it isn't convolution-based like those effects (both of which increase the processing load on your system). As a result, you can make real-time changes more quickly and effectively in Multitrack View, without needing to lock effects to a track.

Although QuickVerb is not convolution-based and is most similar to Studio Reverb, the latter works slightly faster, has better sound quality, and has more options for better control and tonal variation.

To use the Studio Reverb effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Delay Effects, and double-click Studio Reverb.
- 3 Set the desired options.

For more information, see [“Studio Reverb options” on page 169](#).

Studio Reverb options

The Studio Reverb effect provides the following options:

Room Size Sets the room size.

Decay Adjusts the amount of reverberation decay in milliseconds.

Early Reflection Controls the percentage of echoes that first reach the ear, giving a sense of the overall room size. Too high a value can result in an artificial sound, while too low a value can lose the audio cues for the room's size. Half the volume of the original signal is a good starting point

Stereo Width Varies the reverb signal between the stereo channels. A full left setting produces a mono reverb signal.

High Frequency Cut Specifies the highest frequency at which reverb can occur.

Low Frequency Cut Specifies the lowest frequency at which reverb can occur.

Damping Adjusts the amount of attenuation applied to the high frequencies of the reverb signal over time. Higher percentages create more damping for a "warmer" reverb tone.

Diffusion Simulates the absorption of the reverberated signal as it is reflected off of surfaces, such as carpeting and drapes. Lower settings create more echoes, while higher settings produce a smoother reverberation with fewer echoes.

Original Signal (Dry) Sets the percentage of source audio to leave in the effect.

Reverb (Wet) Sets the percentage of reverb to place in the effect.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see ["Working with effects" on page 33](#).

Creating special effects

Effects commands in the Special menu let you introduce processing effects that are both innovative and wild. You can use the Convolution effect to use one waveform to modify another, the Distortion effect to make a waveform sound as if it's coming from an overdriven amplifier or speaker, and the Music effect to create notes and chords sampled from a waveform.

Using the Convolution effect

The Convolution effect multiplies every sample in one wave (the *impulse*) by the samples contained in another waveform. (An "impulse" is the data by which every other sample in a waveform is multiplied. For instance, if the impulse is a single sample of a full volume "click" sound then the convolution of that impulse with any audio data is just that audio data itself. If that click is at half volume, then the convolution is the audio data at half volume.) In a sense, this effect uses one waveform to "model" the sound of another waveform. The result can be that of filtering, echoing, phase shifting, or any combination of these effects. That is, any filtered version of a waveform can be echoed at any delay, any number of times.

For example, modeling someone saying "Hey" with a drum track (short, full-spectrum sounds such as snares work best) results in the drums saying "Hey" each time they are hit. You can build impulses from scratch by specifying how to filter the audio and what delay rate to apply, or by copying audio directly from a waveform.

With the proper impulses, you can simulate any reverberant space. For example, if you have an impulse of your favorite cathedral, and you convolute it with any mono audio (for which the left and right channels are the same), the result sounds as if that audio were played in that cathedral. You can generate such an impulse by going to the cathedral, standing where you want the audio to seem to be coming from, generating a loud impulsive noise (like a “snap” or “click”), and recording the noise in stereo. If you use this recording as an impulse, convolution with it sounds as if the listener is at the exact location of the recording equipment, and the convoluted audio is at the location of the snap or click.

If several ticks descend in amplitude over time, such as one tick every 100 milliseconds, with each tick half as loud as the previous one, then the resulting convolution with audio has 100 milliseconds between each echo, and each echo is half as loud as the previous one.

To get a feel for Convolution, open and play with some of the sample Impulse (.imp) files that come with Adobe Audition. You can find them in the Imps folder within the folder for Adobe Audition and on the Adobe Audition CD.



Use convolution to sustain a sound for any length of time. For example, the sound of a person singing “aaaaaah” for one second can be turned into thousands of people singing “aaaaaah” for any length of time by using dynamically expanded white noise. Also, to send any portion of an unprocessed “dry” signal back out, simply add a full spectrum echo at 0 milliseconds. The Left and Right volume percentages are the resulting volume of the dry signal in the left and right channels.

To use the Convolution effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Special, and double-click Convolution.
- 3 Set the desired options.

For more information, see [“Convolution options” on page 170](#).

Convolution options

The Convolution effect provides the following options:

Load Opens a previously saved impulse (such as the ones that come with Adobe Audition). The sample rate of an impulse affects the outcome of convolution. For example, if an impulse is created at 44100 Hz, and it is later reopened and used on a 22050 Hz file, everything is stretched out 2:1. Filtered echoes are at half the frequency, and delays are twice as long.

Save Saves an impulse.

Clear Clears an impulse completely.

Mono, Stereo Specifies how the impulse works with mono or stereo data. Mono impulses work with either mono or stereo data (the left and right channels are convoluted with the same impulse). Stereo impulses convolute the left and right channels separately.

Scaled By Sets the scaling factor to use when you add a highlighted selection to an impulse to determine its volume. By default, Adobe Audition provides a good starting value. Lower this value to increase the amplitude of the impulse. Note that any audio can be added to an impulse directly.

Minimum Sets the lower cutoff frequency of the echo when bandpassed echoes are added. For example, to echo just the range from 500 Hz to 1000 Hz, enter 500 Hz for the minimum value.

Maximum Specifies the upper cutoff frequency of the echo when bandpassed echoes are added. For example, to echo just the range from 500 Hz to 1000 Hz, enter 1000 Hz for the maximum value.

FIR Size Sets the size of the FIR filter to use to generate the filtered echo.

Note: Adobe Audition recommends a minimum delay (listed below FIR Size) when you add this echo. If you use a smaller delay than suggested, the echo may contain more frequencies than you want. You can ignore this delay for full spectrum echoes, as they are just single sample ticks in the impulse.

Add Sel Adds the current selection to the impulse at the delay and at the left and right volumes specified. You can add as many selections of actual audio as you like.

Note: You can make any audio data part of an impulse by first highlighting the audio and then clicking Add Sel. Ordinarily, you should first scale down any such selection to a lower volume; otherwise, the convolution will be extremely loud.

Add Echo Adds the bandpassed echo to the impulse at the delay and at the left and right volumes specified. You can add as many echoes as you like.

Note: To add a tick at any volume, enter the Left and Right volume percentages and the Delay at which the tick should appear. Doing so creates an echo of the specified volume at the given delay after convolution. Besides just echoes, you can add filtered versions of echoes by entering the minimum and maximum frequencies to echo. To echo all frequencies outside the range, add a full spectrum echo (for example, from 0 Hz to 22,050 Hz) at a specific delay, and then add another echo at the same delay but with different minimum and maximum values, and inverted Left and Right percentages (for example, – 100% instead of 100%).

Delay Sets the number of milliseconds by which samples are delayed. For pre-echoes, place at least one full-spectrum echo (Minimum = 0 Hz, Maximum = 22050 Hz) at a longer delay (such as 1000 milliseconds). Then, any echo placed before 1000 milliseconds is a pre-echo.

Left Specifies the percentage of left volume to add.

Right Specifies the percentage of right volume to add.

Volume Adjusts the volume level if the convoluted audio comes out too soft or too loud.

Shift Compensates for convoluted audio that migrates too far to the right with respect to the original audio. In general, set this value to half the FIR size for impulses you build from scratch to compensate for the delay incurred when the minimum delay used is only half the FIR size.

View Left Displays the impulse for the left channel only.

View Right Displays the impulse for the right channel only.

View Both Displays the impulse for both channels.

Normalized View Displays the impulse's amplitude so that it fits precisely in the graph vertically.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Distortion effect

Use this effect to simulate blown car speakers, muffled microphones, or overdriven amplifiers. Have fun making your audio sound really bad or adding fuzz to guitar licks to get that authentic heavy metal sound.

To use the Distortion effect:

- 1 Select an audio range (Edit View) or track (Multitrack View).
- 2 In the Effects tab of the Organizer window, expand Special, and double-click Distortion.
- 3 Set the desired options.

For more information, see [“Distortion options” on page 172](#).

Distortion options

The Distortion effect provides the following options:

Symmetric Changes the tabs in the dialog box from (Symmetric)/(Symmetric) to Positive/Negative.

(Symmetric)/(Symmetric) tabs Create identical positive and negative curves. You cannot choose one over the other.

Positive/Negative tabs Let you specify separate distortion curves for positive and negative sample values. Select the tab corresponding to the distortion curve you want to display.

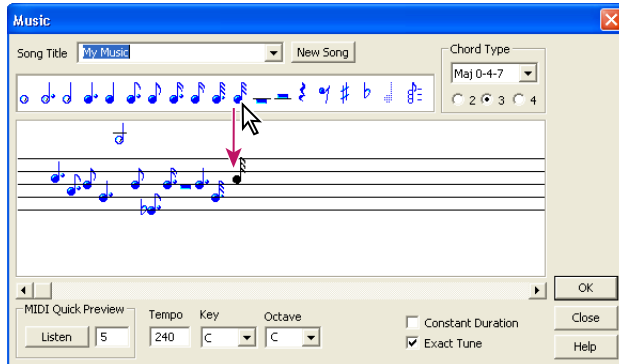
Copy From Positive Copies the positive curve to the Negative window. This option appears on the Negative tab.

Distortion graph Depicts the input sample value in decibels along the horizontal ruler (x-axis) and the output sample value in decibels along the vertical ruler (y-axis). The default line that flows directly from the lower left to the upper right depicts an unmodified signal, since every input value goes to the matching output value. Adjust the shape of this line to adjust the input to output assignments. The readout below the graph displays the current input and output sample values that correspond to the cursor position.

Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Music effect (Edit View only)

The Music effect lets you use any short selection as a “voice” to synthesize music or harmonize a wave using a particular chord. While this function is far from a complete MIDI authoring studio, it provides a quick and simple way to put a sample to music.



The Music dialog box

To use the Music effect:

- 1 In Edit View, select the part of the waveform you want to use as a quarter note.

Note: This selection must be under ten seconds long. If you don't select a range, Adobe Audition uses the data on the clipboard instead. Keep in mind that the clipboard data is filled with the sample automatically after music is generated. Thus, selecting music a second time automatically uses the last sample.

- 2 In the Effects tab of the Organizer window, expand Special, and double-click Music.
- 3 Set the desired options.

For more information, see [“Music options” on page 173](#).

Music options

The Music effect provides the following options:

Song Title Lets you name the song you compose or choose previously saved songs. The song data is saved in the Songs.ini file in the Adobe Audition data folder.

New Song Clears the song title and staff.

Notes, rests, and staff Lets you compose a song that contains up to 256 notes and rests. Simply drag the notes and rests you want to the staff. To sharpen or flatten a note, drag the sharp or flat symbol on top of the note you want to transpose. To clear a sharp, flat, or chord from a note, drag the faded-looking quarter-note onto the note you want to clear. To remove a note, drag it off of the staff. Use the horizontal slider to scroll through a song.

Chord Type Establishes the major or minor chord type and voicing to use. Select 2, 3, or 4 voices, and choose a chord type from the list. Then, drag the chord object (which looks like three stacked notes) onto a note in the staff. The note you drop it on becomes the starting note of the chord, and the other notes automatically appear above it in the right ratios.

Listen Previews the sequence if you have MIDI playback capabilities. Playback begins at the leftmost note visible on the staff and continues to the end of the song. The song plays through channels 1 and 13 for Extended and Base level compatibility. Set the desired instrument by entering its MIDI instrument number in the text box to the right.

Tempo Specifies the tempo at which to play the sequence, measured in quarter notes (beats) per minute. The sample's length is the length of a quarter note. If a note is longer than the period determined by the tempo, the notes will overlap.

Key Specifies the key for the song. Only standard major key signatures are listed, so to specify a minor key, choose the relative major (B flat for C# minor, for example).

Octave Lets you transpose a sequence by octaves. Choose "C" to play notes at normal transposition.

Constant Duration Makes all notes the same length as the original sample, regardless of pitch. This operation takes extra time to process, but high-pitched notes will be the same length as lower-pitched notes. If this option is unselected, the note is created by directly stretching or compressing the original sample, resulting in higher pitches being shorter than lower pitches.

Exact Tune Tunes the sample so that when played at A (above middle C), its frequency is 440 Hz. If this option is unselected, the sample's original frequency is played at A (above middle C).

Using multitrack-only effects


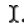
Some effects in Adobe Audition are available only in Multitrack View. The Effects menu and the Multitrack category in the Effects tab contain all of these effects. For information about selecting clips and ranges in Multitrack View, see ["Selecting ranges in the track display" on page 182](#).

Using the Envelope Follower effect (Multitrack View only)

The Envelope Follower effect varies the output level of one waveform based on the input level of another. The amplitude map, or envelope, of one waveform (the analysis wave) is applied to the material of a second waveform (the process wave), resulting in the second waveform taking on the amplitude characteristics of the first. This effect lets you, for example, have a bass guitar line that sounds only when a drum is hit. In this example, the drum waveform is the Analysis wave, and the bass guitar waveform is the Process wave.

In addition to applying an amplitude envelope to a waveform, you can alter the dynamic properties of the resulting signal with a variety of settings to otherwise expand, gate, compress, or limit it.

To use the Envelope Follower effect:

- 1 In Multitrack View, position the wave clips so that the sections you want to process together are aligned.
- 2 Select the Hybrid tool  or the Time Selection tool .
- 3 In the track display, select the range you want to process.
- 4 Ctrl-click the wave clips you want to process.

Note: If you select a range by dragging across a clip, that clip is selected by default; if you Ctrl-click the clip, you deselect it.

- 5 In the Effects tab of the Organizer window, expand Multitrack, and double-click Envelope Follower.

- 6 Set the desired options.

For more information, see ["Envelope Follower options" on page 175](#).

Envelope Follower options

The Envelope Follower effect provides the following options:

Analysis Wave Specifies which waveform to use as the analysis wave. In other words, it provides the amplitude envelope. You can choose any waveform in the session window.

Process Wave Specifies which waveform to apply the analysis waveform to. You can choose any waveform in the session window.

Output To Specifies which track to output the resulting waveform to. The default setting is the next available track.

Gain Processor Provides the following options:

- Output Gain specifies the amount of gain (measured in decibels) added to the output signal.
- Attack Time determines how long the processed output signal takes to reach the specified output volume. For example, if a portion suddenly drops 30 dB, it takes time specified before the output drops to its corresponding volume level. If the sum of the Attack and Release times is too short (less than 20 milliseconds total), audible effects, such as a vibrating sound, occur at around 1000 Hz/milliseconds total. For example, if the Attack and Release times are each set to 5 milliseconds (10 milliseconds total), then a vibrating sound occurs at 100 Hz. A total of 30 milliseconds is about as low as you can go without introducing these effects.
- Release Time determines how long the previous output level takes to reach the specified output volume. For example, where the Attack Time is how long the start of a pulse takes to reach the desired output volume, the Release Time is how long the end of the pulse takes to reach the desired level.
- Joint Channels uses both channels to find a single input decibel value and amplifies both channels by the same amount, preserving the stereo center-channel image. With stereo files, each channel can be compressed independently, sometimes causing the surrounding background noise to get louder on one channel. For example, a loud drumbeat in the left channel makes the background noise in the right channel louder than in the left.

Level Detector Provides the following options:

- Input Gain specifies the amount of gain (measured in decibels) added to the signal before it goes into the Level Detector (the section that detects the current level). The input gain essentially moves the graph plot up or down.
- Attack Time determines how long the processed output signal takes to reach the specified output volume. For example, if a portion suddenly drops 30 dB, it takes time specified before the output drops to its corresponding volume level. If the sum of the Attack and Release times is too short (less than 20 milliseconds total), audible effects, such as a vibrating sound, occur at around 1000 Hz/milliseconds total. For example, if the Attack and Release times are each set to 5 milliseconds (10 milliseconds total), then a vibrating sound occurs at 100 Hz. A total of 30 milliseconds is about as low as you can go without introducing these effects.
- Release Time determines how long the previous output level takes to reach the specified output volume. For example, where the Attack Time is how long the start of a pulse takes to reach the desired output volume, the Release Time is how long the end of the pulse takes to reach the desired level.

- Peak mode is provided for backward compatibility. It is a graph interpretation method that is slightly outdated and a bit more difficult to use than RMS. It is equivalent to twice the RMS value (for example, –20 dB in RMS mode equals –40 dB in Peak mode).
- RMS is a graph interpretation method that more closely matches the way people hear volume. This mode causes the output to be exactly the RMS amplitude that is specified in the graph. For example, a limiter (flat horizontal line) at –10 dB causes the RMS amplitude of the result to average –10 dB (where 0 dB is a maximum amplitude sine wave without clipping).

Low Cutoff Specifies the lowest frequency that dynamics processing affects.

High Cutoff Specifies the highest frequency that dynamics processing affects.

Lookahead Time Helps to handle sharp spikes that might occur at the onset of a louder signal by starting the attack the specified number of milliseconds before the audio becomes loud, instead of right on top of the transient. The spikes occur because of the time needed to determine and react to the current signal level (as determined by the attack values). For brief moments, these transients can go beyond the limits of the compressor settings. You might want these spikes in certain compression scenarios to enhance the impact of, say, a drum hit, but they aren't desirable if you're using limiting to reduce the maximum amplitude. Otherwise, with Lookahead Time set to 0, a spike stays loud until all of the attack times elapse.

Graph Depicts input levels along the horizontal ruler (x-axis) and the new output level along the vertical ruler (y-axis). An unedited signal is depicted as a straight line from the lower left to the upper right of the graph, since every input value goes to the matching output value. Adjusting the line adjusts the input and output assignments, thereby altering the dynamic range. For example, you can boost all input that has a level of around –20 dB, leaving everything else unchanged.

Flat Resets the graph to its default state.

Invert Inverts the graph. You can invert a graph only if it has points in the two default corners (–100, –100 and 0, 0) and if its output level increases from left to right (that is, each control point must be higher than the one to its left).



Note: For information on graphs, spline curves, presets, previews, and other options that are common to many effects dialog boxes, see [“Working with effects” on page 33](#).

Using the Frequency Band Splitter effect (Multitrack View only)

The Frequency Band Splitter lets you take a selected waveform clip (or a highlighted section thereof) and make up to eight copies of it, with each copy assuming a different frequency range of the original. Split points are determined by the crossover frequencies you specify. Each copy of the waveform is placed in its own track in the session window. You can then edit or apply effects to each band separately.

For example, using the default setting of three bands with crossover values of 800 and 3200 creates three copies of the selected waveform: one with the frequencies of the selected wave from 0 to 800 Hz, one from 800 to 3200 Hz, and one from 3200 to 22050 Hz (or whatever the maximum frequency present is, based on the sample rate).

To use the Frequency Band Splitter effect:

1 In the track display in Multitrack View, select the clip or the range you want to process. (Use the Hybrid tool  or the Time Selection tool  to select a range.)

2 Ctrl-click the wave clip.

Note: If more than one wave clip is selected, the Frequency Band Splitter effect is unavailable. Also, if you select a range by dragging across a clip, that clip is selected by default; if you Ctrl-click the clip, you will deselect it.

3 In the Effects tab of the Organizer window, expand Multitrack, and double-click Frequency Band Splitter.

4 Set the desired options.

For more information, see [“Frequency Band Splitter options” on page 177](#).

Frequency Band Splitter options

The Frequency Band Splitter effect provides the following options:

Bands Sets the number of split points. The original waveform is copied the number of times you specify, with each copy having a different frequency range, as determined by the number of crossovers.

Crossovers Determines the crossover frequencies that are used for split points.

Output Waves Specifies the name of each new wave clip. By default, each new clip is assigned the original name of the waveform plus its frequency range.

Max FIR Filter Size Sets the maximum size of the FIR (Finite Impulse Response) filter, which maintains phase errors over the response curve. FIR filters are unlike IIR filters, which can have phase error (often audible as a ringing quality). Higher values create higher accuracy in the frequency filtering. The default value, 320, works well most of the time, but you should increase it if distortion or ringing occurs in the filtered waves.

Note: For information on presets, see [“Working with effects” on page 33](#).

Using the Vocoder effect (Multitrack View only)

A vocoder takes two inputs, usually an instrument and a voice, and modulates one signal (the process signal, usually the instrument) with the other (the control signal, usually the voice). This modulation allows one signal to “control” the other. In the example here, the instrument (the process signal) could be made to “sing” by affecting it with the voice (the control signal).

To use the Vocoder effect:

1 In Multitrack View, position the wave clips so that the sections you want to process together are aligned.

2 Select the Hybrid tool  or the Time Selection tool .

3 In the track display, select the range you want to process.

4 Ctrl-click the wave clips you want to process.

Note: If you select a range by dragging across a clip, that clip is selected by default; if you Ctrl-click the clip, you will deselect it.

5 In the Effects tab of the Organizer window, expand Multitrack, and double-click Vocoder.

6 Set the desired options.

For more information, see [“Vocoder options” on page 178](#).

Vocoder options

The Vocoder effect provides the following options:

Control Wave (Voice) Sets any active waveform in the session, usually a vocal, as the control signal.

Process Wave (Synth) Sets any active waveform in the session as the process signal. This waveform is usually a synthesized sound to replace a vocal.

Output To Sets which track to output the resulting waveform to.

FFT Size Specifies the size of the FFT (Fast Fourier Transform). Higher sizes usually create more precise results, but they take longer to process.

Overlays Sets the number of FFTs that overlap. More overlays can produce smoother results, but they take longer to process. Values of 3 to 12 work well.

Interval Size Lets you specify the number of milliseconds per FFT. Values between 10 and 30 usually work best, unless higher overlay settings are used. Smaller values can produce a hum, while larger values can produce blocky sounding results.

Window Width Lets you specify the percentage per FFT. A value of 90% generally produces good results.

Vocal Crossover Determines the frequency used to filter out, or separate, the source waveform's underlying base frequency (voice) from the vocal formants (the vowel sounds). With higher values, more formants and less of the source voice is carried over. Ideally, you want none of the source voice carried over and all of the formant information, so the synthesizer “talks.”

Resynthesis Window Specifies the width of the window used for resynthesizing the vocoded signal. Narrower windows make hard consonants sound clearer, and with higher overlay settings, they give better time resolution if the vocoded signal sounds too smoothed out. The numbers available for this option are always lower than the number of overlays.

Affect Level Sets the amount of vocoded signal that ends up in the resulting waveform. For example, set this value to 100% for full vocoding, 50% to keep more of the original waveform, and 15% to produce a subliminal effect that barely affects the process wave with the voice.

Amplification Specifies the amplification by which to adjust the final waveform. This value can be zero, but if the results are too quiet or loud, raise or lower this value as needed.



For ease of use, set Window Width to about 90%, use 3 or 4 overlays, set Resynthesis Window to 1 or 2, and choose an FFT size between 2048 and 6400.

Mixing Multitrack Sessions

About mixing multitrack sessions

In Multitrack View, you can add audio, video, ReWire, and MIDI files to separate tracks of a multitrack session and then mix those tracks together. When you're happy with a mix, you can export a mixdown file for use on CD, the Web, and more.

Multitrack View is a flexible editing environment because mixing occurs in real time and is nondestructive. Because mixing occurs in real time, you can change mix settings during playback and immediately hear the results. For example, you can adjust a track's volume as a session plays to properly blend the track with other tracks. Because mixing is nondestructive, mixing adjustments don't permanently change original source files. For example, you can apply four effects to a track and later remove two effects to create a different sonic texture.

Adobe Audition saves information about mix settings and source files in session (.ses) files. Session files are relatively small because they contain only pathnames to source files and references to mix parameters (such as volume, pan, and effect settings). To more easily manage session files, save them in a unique folder with the source files they reference. If you later need to move the session to another computer, you can simply move the unique session folder. For more information, see ["Saving sessions" on page 238](#).

Note: Only one session can be open at a time.

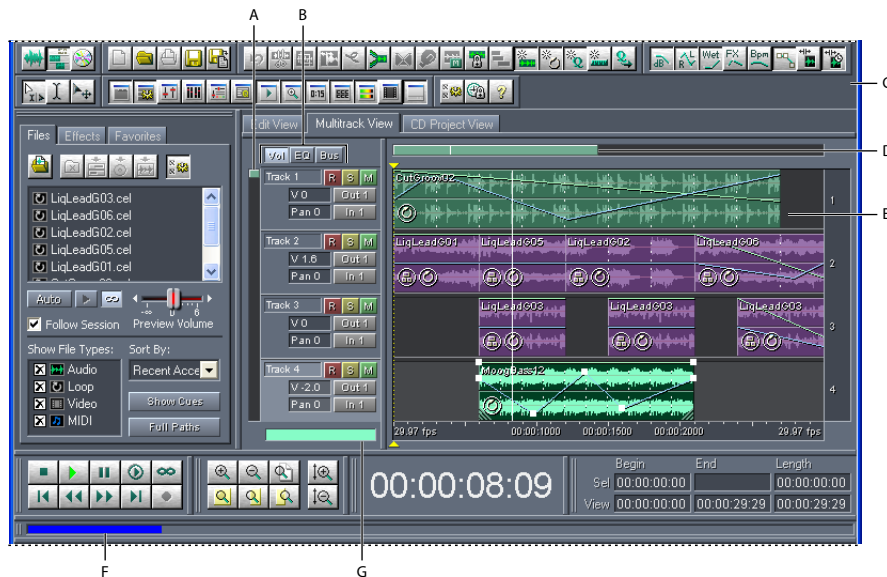


Working with sessions

The Multitrack View work area includes several unique elements that help you mix sessions. On the left, the track controls let you adjust track-specific settings, such as volume and pan. (See [“Working with audio tracks” on page 197](#).) On the right, the track display lets you edit the clips in each track. (See [“Working with clips” on page 187](#).)



For information about elements of the work area that Multitrack View shares with Edit View, see [“About the work area” on page 19](#).



Multitrack View work area

A. Vertical scroll bar **B.** Track controls **C.** Toolbar **D.** Horizontal scroll bar **E.** Track
F. Load Meter **G.** Mix Gauge

Creating new sessions

When you create a new session, you specify its sample rate. (See [“About sample rates” on page 93](#).) Any files added to the session must share this sample rate. If you try to import a file with a different sample rate, Adobe Audition lets you convert it.

You can base new sessions on the default session, borrowing default settings such as device assignments and master volume levels. (See [“Setting the default session” on page 181](#).)



To convert the sample rate of an existing session, use the Save Session As command and save converted copies of all referenced files. (See [“Saving sessions” on page 238](#).)

To create a new session:

- 1 Choose File > New Session.
- 2 Select the desired sample rate.
- 3 If you want to base the new session on the default session, select Use Default Session. (This option appears only if you’ve set a default session.)
- 4 Click OK.

Setting the default session

After you set a default session, it opens when you start Adobe Audition. The default can also serve as a template for new sessions, letting you share settings such as device assignments and master volume levels across multiple sessions.

To set the current session as the default:

- 1 Choose File > Default Session > Set Current Session As Default.
- 2 If the session contains clips, click Yes.



To create a new session that uses the default session as a template, see [“Creating new sessions” on page 180](#).

To not use a default session:

Choose File > Default Session > Clear Default Session.

Inserting or deleting time in a session

You can use the Insert/Delete Time command to insert silence into a session or to delete a selected range from the session.

To insert or delete time in a session:

- 1 Place the current-time indicator at the desired insertion point, or select the range you want to delete.
- 2 Choose Edit > Insert/Delete Time, and set the following options:



Insert Shifts all material (clips or parts of clips) to the right of the current-time indicator by the amount you specify in the text box. Clips are split if necessary, and the specified amount of silence is inserted.

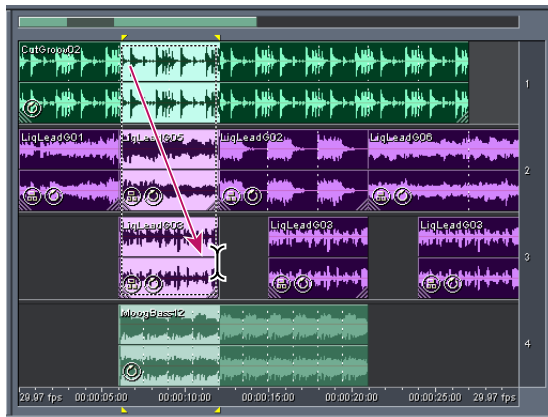
Delete Selected Time Removes the highlighted area and shifts all clips to the right of the selected region.



You must unlock any locked tracks to insert or delete time in a session. To relock such tracks, click the Lock button in the track controls.



Selecting ranges in the track display

To select ranges in Multitrack View, you can use either the Time Selection tool  or the Hybrid tool . Both tools let you select ranges and clips, but the Hybrid tool also lets you move clips. If you prefer to select ranges separately from moving clips, use the Time Selection and the Move/Copy Clip tools rather than the Hybrid tool. (See [“Working with clips” on page 187.](#))



Simultaneously selecting a range and clips in the track display
(Upper three clips are selected, fourth isn't)

To select a range in the track display:

- 1 In the toolbar, select either the Hybrid tool  or the Time Selection tool .
- 2 In the track display, do one of the following:
 - To select only a range, click an empty area of the track display, and drag left or right.
 - To select a range and clips, click a clip, and drag left or right while dragging up or down.


Measuring performance with the Mix Gauge and Load Meter

In Multitrack View, the Mix Gauge and Load Meter help you measure and optimize performance. The Mix Gauge displays the progress of background mixing, a process that Adobe Audition completes whenever you edit a mix (for example, by moving a clip or changing track volume). Background mixing lets you monitor an updated mixdown of a session and is complete when the Mix Gauge reaches 100%. You needn't wait for the Mix Gauge before clicking the Play button, though audio may skip or drop out.

The Load Meter shows the percentage of available CPU power, a particularly important value if you use real-time effects. Unlike the Mix Gauge, the Load Meter indicates a problem if it reaches 100%. At that level, your system will perform erratically because the CPU has no additional processing power. You can reduce CPU load by locking real-time effects. (See [“Locking tracks with real-time effects” on page 207.](#))



Multitrack performance depends primarily upon overall system speed, including CPU and hard disk speed. You can optimize multitrack performance on any system, however, by properly configuring multitrack options in the Settings dialog box. (See [“Multitrack options” on page 51.](#))



Performance indicators in Multitrack View

A. Mix Gauge **B.** Load Meter.

To change background mixing settings:

Right-click the Mix Gauge, and choose any of the following:

- Disable Background Mixing.
- Lower Mix Priority When In Other Applications.
- A Mix Ahead setting to determine how far ahead of the current time Adobe Audition begins mixing. Longer settings allow for faster mix editing, but they might cause drop outs.
- Mix Entire Session to create a new background mix each time you edit a mix.
- A Mix Priority setting to determine the processing priority of background mixing versus other tasks.



Disabling background mixing can improve performance when you need to extensively edit a mix.

To manually start background mixing:

Choose Edit > Refresh Now.

To view or hide the Load Meter:

Choose Window > Load Meter.

Using sessions as SMPTE masters or slaves

By using sessions as SMPTE masters or slaves, you can synchronize the transport controls of Multitrack View with a MIDI sequencing application or an external hardware device, such as a videotape machine. Before using a session as a master or slave, you must set general SMPTE options that apply to all multitrack sessions. (See [“Setting up for SMPTE synchronization” on page 40.](#))

As a SMPTE master, a session generates timecode in the SMPTE time format you select for the timeline. As a SMPTE slave, a session receives timecode generated elsewhere, reporting the following synchronization statuses in the left of the status bar:

- Opened MIDI Input Device when waiting for incoming timecode.
- Synchronizing when establishing synchronization. (Adobe Audition requires about 5 seconds of timecode, known as *preroll*, to establish synchronization.)
- Playback Synchronized when synchronization is established.

Note: Adobe Audition sends and receives timecode through the MIDI Out and MIDI In ports of your system. To configure these ports, see [“Setting up for SMPTE synchronization” on page 40.](#)

To use a session as a SMPTE master:

- 1 Choose Options > SMPTE Master Enable.
- 2 Select the desired SMPTE time format for the timeline (see [“Monitoring time” on page 63](#)).

To use a session as a SMPTE slave:

- 1 Choose Options > SMPTE Start Offset, click Format, and select the desired SMPTE time format.
- 2 Enter the desired start point in the SMPTE Start Time Offset box, and then click OK. (This option defines Adobe Audition’s start point; it doesn’t offset incoming timecode.)

Note: If you chose the SMPTE Drop time format, the offset must compensate for dropped frames. For example, you must enter 1:00:02 to achieve an offset of 1:00:00.

- 3 Choose Options > SMPTE Slave Enable.

Setting advanced session properties

In the Advanced Session Properties dialog box, you can adjust session-specific mixing, tempo, and metronome settings. You can also add session notes, which can help you recall details about a session or communicate those details to someone else.



To set loop-related session properties, use the Session Properties window. See [“Setting the tempo, time signature, and key for sessions” on page 219](#).

To set advanced session properties:

- 1 Choose View > Advanced Session Properties.
- 2 Set options as desired, and click OK.

General options for sessions

In the General tab of the Advanced Session Properties dialog box, you can set the following options:

SMPTE Start Time Offset Specifies a time offset, which is the time location at which Adobe Audition will begin playback. Click Format to choose a time format. For more information about time formats, see [“Monitoring time” on page 63](#).

Key for Voiced Loops Specifies the session’s key. Any loops that have a key specified will be pitch-shifted to match the session’s key.

Mixing options for sessions

In the Mixing tab of the Advanced Session Properties dialog box, you can set the following options:

Pre-Mixing Specifies the bit depth used for the background mixing process. Best quality is achieved by choosing the 32-bit setting, especially for single output device situations or if you have a fast hard drive. However, if you're using multiple sound cards, it may be advantageous to pick 16-bit for pre-mixing, as less data will be transferred across the hard drive(s), which will speed things up.

Note: *Even if you do 16-bit pre-mixing, no data is lost, and a final mixdown can still be 32-bit.*

Panning Sets the method used for panning waveforms in a multitrack session.

- L/R Cut Logarithmic (default) achieves left-panning by reducing the volume of the right channel, and right-panning by reducing the left channel volume. The channel being panned to doesn't increase in volume as panning gets closer to 100%.
- Equal-Power Sinusoidal pans left and right channels with equal power, so a hard pan contains the same perceived loudness as a center pan. For example, a hard pan to the left amplifies the left channel by 3 dB. This option reflects analog mixing boards.

Note: *Because panning can actually make one channel louder than the original waveform, audible clipping can occur in 16-bit sessions. To avoid clipping, work in the 32-bit realm if you use the Equal-Power Sinusoidal method.*

Master Vol Determines the order of the FX bus and Master Volume slider. Select one of the following:

- Pre-Bus FX to place the FX bus prior to the Master Volume slider. Use this option if you plan to master an exported mixdown file in Edit View.
- Post-Bus FX to place the FX bus after the Master Volume slider. Use this option to master as you mix in Multitrack View.

Volume Envelopes Determines the range of volume envelopes. Select one of the following:

- 0% to 100% Range to place 100% volume at the top of clips.
- 0% to 200% Range to place 100% volume at the center of clips.

Set As Default Stores any changes to mixing options as the default settings.



To set default mixing, playback, and recording options for sessions, use the Multitrack tab in the Settings dialog box. (See [“Multitrack options” on page 51.](#))

Tempo options for sessions

In the Tempo tab of the Advanced Session Properties dialog box, you can set the following options:

Tempo Specifies session tempo. Any loops that have a tempo specified will be time-stretched to match the session's tempo. The time-stretch method used is determined by the individual loop's settings. If you need more detailed tempo settings, right-click on the horizontal ruler in Adobe Audition's Session Display.

- Beats/Minute specifies tempo.
- Beats/Bar specifies time signature.
- Beat/Length specifies the length of each beat. (Enter 4 for a quarter note, 8 for an eighth note, and so on.)
- Ticks/Beat specifies the number of ticks per beat.

Offset Choose from the following:

- Cursor At shows the current position of the cursor in bars-and-beats format. You can also enter a new cursor position here and Adobe Audition will recalibrate the session, including negative time. This is so you can line up your metronome with a song that's already been created and you're adding more to. Any point can be at any time, so that you know the metronome beat will be in the right place.
- Reset 1:1 to Cursor sets the cursor to the first beat of the first measure.
- Song Start shows the exact position (stated in milliseconds) of the beginning of the song.

Metronome options for sessions

In the Metronome tab of the Advanced Session Properties dialog box, you can set the following options:



To directly access metronome settings, choose Options > Metronome.

Enable Metronome Lets you keep perfect time by using the built-in metronome of Adobe Audition. Check the Enable Metronome box to hear the ticks of the metronome whenever you play or record in Multitrack View. (The metronome sound effects won't be recorded—unless, of course, you use a microphone and the computer's speakers are on.)

Sound Set Determines the type of metronome sound.

Volume Sets metronome volume. The default is – 6 dB. A higher number (such as – 3 dB) makes the metronome louder, while a lower number (like – 10 dB) decreases the metronome volume.

Signature Sets the metronome time signature. The current signature pattern is displayed in the noneditable Pattern field. Custom time signatures may be added using the Add Custom button.

Add Custom Opens the Customize Metronome Time Signature dialog box, where you can add customized options to the Signature menu. Set the following options, and then click Add:

- Name determines the name that appears in the Signature menu.
- Pattern specifies the pattern of the time signature. Use a 0 for no beat, a 1 for a down beat, a 2 for a secondary beat, and 3 for a regular beat. Spaces between the numbers are optional.
- Beats/Bar specifies tempo.
- Beat Length specifies the length of each beat. (Enter 4 for a quarter note, 8 for an eighth note, and so on.)

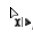


Pattern Shows the current metronome pattern.


Notes for sessions

In the Notes tab of the Advanced Session Properties dialog box, you can type session notes, which can help you recall details about a session or communicate those details to someone else. For example, you could note the musicians for the session, indicate the instrument each played, and detail alternative EQ settings for those instruments.

Working with clips

When you insert an audio, MIDI, or video file in Multitrack View, the file becomes a clip on the selected track. You can easily move clips to different tracks or timeline positions. You can also edit clips nondestructively, trimming their start and end points, crossfading them with other clips, and more.

To work with clips in the track display, you can use either the Hybrid tool , which lets you move clips and select ranges, or the Move/Copy Clip  and Time Selection  tools, which separate these tasks. (See [“Selecting ranges in the track display” on page 182.](#))



Aligning and grouping two clips

Selecting and moving clips

To move a clip or change its properties, you must select it. You can select either individual clips or all clips in a track or session.

To select an individual clip:

Click the clip in the track display.

To select all clips in a track:

- 1 Select the track.
- 2 Choose Edit > Select All Clips In Track [number].




If space exists between clips, double-click that space to quickly select all clips in a track.

To select all clips in a session:

Choose Edit > Select All Clips.


To move selected clips:

- 1 Select the Move/Copy Clip tool  in the toolbar.
- 2 Drag the clips.



If you prefer, select the Hybrid tool, and then right-click and drag the clips.

Grouping clips

You can group clips to more efficiently organize, edit, and mix a session. For example, you can group guitar clips together to easily identify, select, and move them. Grouped clips appear with the group icon  and in a different color than ungrouped clips.



Changes to clip mute and lock properties affect all audio clips in a group. See [“Setting audio clip properties” on page 193](#).

To group clips:

- 1 Hold down the Ctrl key, and click each clip you want in the group.
- 2 Choose Edit > Group Clips. Alternatively, right-click any clip in the group, and choose Group Clips.

To ungroup clips:

Select any clip in the group, and choose Edit > Group Clips. Alternatively, right-click any clip in the group, and deselect Group Clips.

To change the color for a group:

- 1 Select any clip in the group, and choose Edit > Group Color. Alternatively, right-click any clip in the group, and choose Group Color.
- 2 Select a color, and click OK.

Aligning clips

You can align the left or right edges of multiple clips, giving them the same start or end point.

To align clips:

- 1 Hold down Ctrl, and select the clips.
- 2 Choose Edit > Align Left or Edit > Align Right.

Note: Because the relative position of grouped clips is fixed, you must ungroup them to align them.

Snapping clips to loop endpoints and other clips

Snapping lets you quickly align clips with loops and other clips. If snapping is enabled, both dragged clips and the current-time indicator snap to loop endpoints and clip edges.

💡 The procedure in this section describes snapping options that are unique to Multitrack View. For information about snapping options that Multitrack View shares with Edit View, see [“Snapping” on page 78](#).

To set snapping options for clips:

Choose Edit > Snapping, and choose from the following options:

Snap To Clips Causes clips to snap to the beginning or end of other clips.


Snap To Loop Endpoints Causes clips to snap to the beginning or end of loops.

💡 While you drag a clip, a white line appears in the track display when snapping points meet. For example, if Snap To Clips is selected, the white line appears when a clip is aligned with the beginning or end of another clip.

Editing audio and MIDI clips

You can edit audio and MIDI clips to suit the needs of a mix. Because Multitrack View is nondestructive, clip edits are impermanent; you can return to the original, unedited clip at any time. If you want to permanently edit an audio clip, however, you can quickly open the source file in Edit View.

You can edit audio and MIDI clips in many different ways. After selecting a range of a clip, you can cut out that range or trim the clip to it. You can adjust the edited boundaries of a clip, revealing or hiding more of it. You can also slip edit a clip to move its contents but not its boundaries.








Using the Adjust Boundaries command to reveal more of a previously edited clip






Though the procedures in this section mention toolbar buttons for clip editing commands, you can also access these commands from the Edit menu or the clip context menu. For example, you can choose Edit > Trim instead of clicking the Trim To Selection button.



To edit a clip with a selected range:

- 1 In the toolbar, click the Time Selection tool  or the Hybrid tool .
- 2 Drag across the clip to select both it and a range.
- 3 In the toolbar, do one of the following:
 - To trim the clip to the range, click the Trim To Selection button .
 - To cut the range from the clip, click the Cut Wave(s) Out Of Selection button . (Alternatively, press Delete.)
 - To adjust clip edges to the range, click the Adjust Waveform Boundaries To Selection button . (To reveal more of a previously edited clip, extend the range beyond the current clip edges.)

To edit clip edges by dragging:

- 1 In the toolbar, click the Clip Edge Dragging button .
- 2 In the track display, position the cursor over the left or right edge of the clip. The edge-dragging icon  appears. (If instead the time stretch icon  appears, position the cursor above the corner handle.)
- 3 Drag to edit clip edges.

To slip edit a trimmed or looped clip:

- 1 In the toolbar, click the Move/Copy Clip tool  or the Hybrid tool .
- 2 Hold down Alt, and right-click drag across the clip.

To return to the full, original version of a clip:

Select the clip, and choose Edit > Full. Alternatively, right-click the clip, and choose Full.




The Full command doesn't apply to loops; instead, adjust clip boundaries by dragging them.

To edit the source file for an audio clip in Edit View:

Double-click the clip.

Splitting and rejoining audio and MIDI clips

The Split command functions similarly to a traditional tape splice; it cuts audio and MIDI clips into parts. When a clip is split, each part becomes a new clip that can be independently moved or deleted. Splitting is nondestructive, so you can rejoin split clips with the Merge/Rejoin Split command.



Selecting a range and splitting one clip into three independent clips

To split a clip:

- 1 In the toolbar, click the Time Selection tool or the Hybrid tool .
- 2 Do either of the following:
 - To split the clip in two, click where you want the split to occur.
 - To split the clip into three, drag across it to specify two split points (one at the beginning of the selection; one at the end).
- 3 In the toolbar, click the Split Clip button .

To rejoin split clips:

- 1 In the toolbar, click the Move/Copy Clip tool or the Hybrid tool .
- 2 Position the clips beside each other on the same track.
- 3 Right-click one of the clips, and choose Merge/Rejoin Split.


Copying audio and MIDI clips

You can create two types of copied audio clips: reference copies that share source files and unique copies that have independent source files. You can create only reference copies of MIDI clips. For audio clips, the type of copy you choose depends upon the amount of available disk space and the nature of destructive editing you plan to perform in Edit View.

Reference copies consume no additional disk space, letting you simultaneously edit all instances by editing the original source file. (For example, you can add the Flanger effect to the source file in Edit View and automatically apply the effect to all 30 referenced copies in a session.)

Unique copies have a separate audio file on disk, allowing for separate editing of each version in Edit View. (For example, you can add destructive effects to the version in an introduction while leaving the version in a verse dry).

To copy a clip:

- 1 Click the Move/Copy Clip tool  in the toolbar.
- 2 Right-click and drag the clip.
- 3 Release the mouse button, and choose either of the following from the pop-up menu:
 - Copy Reference Here
 - Copy Unique Here



If you prefer, copy clips with the Hybrid tool. To copy a reference clip, hold down Shift and right-click drag. To copy a unique clip, hold down Ctrl and right-click drag.

Repeating audio and MIDI clips

With the Clip Duplicate command, you can duplicate repetitions of a clip in a track without consuming additional disk space. You can also specify the spacing between each repetition.



To copy a clip to a different track or to irregular positions in the current track, see [“Copying audio and MIDI clips” on page 191](#).

To repeat a clip:

- 1 Select the clip, and choose Edit > Clip Duplicate.
- 2 Set the following options:

Duplicate Clip Specifies the number of times to duplicate the clip.

Spacing Determines the spacing between each duplicated clip:

- No Gaps—Continuous Looping places each duplicate directly after its preceding clip, for a continuous loop.



For a more flexible method of looping, adjust a clip’s loop properties. See [“About loops” on page 214](#).

- Evenly Spaced defines the spacing between each clip according to the time display format. This value defaults to the length of the selected clip, producing the same effect as the No Gaps option. Enter a greater value to place space between each clip, or enter a lesser value to overlap clips.



To repeat a clip such as a drum hit at every other beat in a song, set the time format to Bars And Beats. (See [“Monitoring time” on page 63](#).) If the clip’s start and end points don’t align properly with beats, trim the clip in Edit View by using Edit > Find Beats.

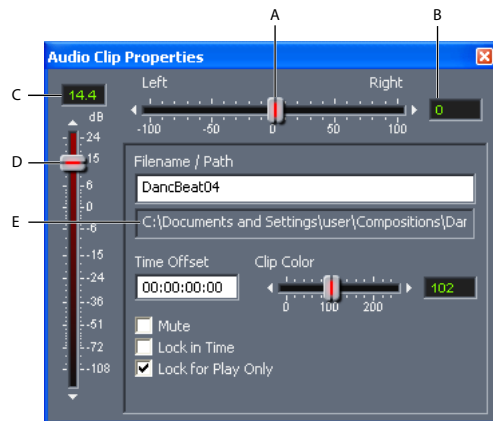
Setting audio clip properties

In the Audio Clip Properties window, you can change settings such as volume, pan, and color for audio clips. Clip settings for volume, pan, and mute are independent from similar track controls.

You can also lock clips in time and lock them for play only. If a clip is locked in time, you can move it up or down to another track, but you can't move it right or left to a new timeline position. If a clip is locked for play only, you can record in the remainder of the track without recording over the clip.




You can directly access many audio clip properties from the Edit menu or the clip context menu (for example, choose Edit > Adjust Audio Clip Volume).



Audio Clip Properties window

A. Pan slider **B.** Pan text box **C.** Volume text box
D. Volume slider **E.** Pathname for source file

To change the properties of an audio clip:


- 1 Right-click the clip, and choose Audio Clip Properties.
- 2 Do any of the following:
 - To change volume, pan, or color, drag the volume, pan, or color slider to the desired position.
 - To lock the clip in time, select Lock In Time. A lock icon  appears on the clip.
 - To lock the clip for play only, select Lock For Play Only. If the containing track is record-enabled, the clip remains the same color; other clips in the track turn red.
 - To mute the clip, select Mute.
 - To move the clip to a new timeline position, enter a start time in the Time Offset text box.
 - To change the clip name, type in the Filename text box. (When you save the session, Adobe Audition prompts you to save a copy of the source file with the new clip name.)

Crossfading audio clips

You can crossfade audio clips to transition smoothly from the end of one clip to the beginning of another. Crossfades consist of a fade out and a fade in over a transition region. To create a smooth transition, select a transition region that starts before the end of the first clip and extends beyond the beginning of the second clip.



The fade curves created with Crossfade commands are volume envelopes, which you can edit. See [“Automating mixes with clip envelopes” on page 207](#).



Selecting a range and two clips, and applying a linear crossfade

To crossfade two clips:

- 1 Place the clips on separate tracks.
- 2 Position the clips so the end point of the first overlaps the start point of the second.
- 3 Across the overlapping area, select a transition region for the crossfade.



To precisely place the start and end points for the crossfade at clip start and end points, choose Edit > Snapping > Snap To Clips.


- 4 Ctrl-click both clips.
- 5 Choose Edit > Crossfade, and then choose one of the following:
 - Linear to produce an even crossfade.
 - Sinusoidal to produce a crossfade with a curved, sine-like slope.
 - Logarithmic In to fade in logarithmically, producing a steeper slope at the end of the fade.
 - Logarithmic Out to fade out logarithmically, producing a steeper slope at the beginning of the fade.

Time stretching audio clips

Time stretching lets you change the length of an audio clip without changing its pitch. This technique is particularly helpful for fitting audio clips to video scenes or layering clips for sound design. You can quickly time stretch a clip either by dragging or setting time stretch properties. When you time stretch by dragging, Adobe Audition analyzes a clip's contents and attempts to select the most natural sounding time-stretch method. When you set properties for time stretching, you also specify which method of time stretching to use.


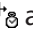
Like other features in Multitrack View, time stretching is nondestructive, so you can disable it at any time.

Note: Time stretching changes the tempo of a clip. If you time stretch a loop-enabled clip, it won't match the session tempo.



Dragging to time stretch a clip

To time stretch a clip by dragging:

- 1 In the toolbar, click the Clip Time Stretching button .
- 2 Select the clip, and then position the cursor over the clip's bottom left or right handle—the time stretch icon  appears.
- 3 Drag the handle to lengthen or shorten the clip.



To temporarily enter time stretching mode, hold down Ctrl, and drag a clip handle.

To set time stretch properties:

- 1 Right-click the clip, and choose Clip Time Stretch Properties.
- 2 Select Enable Time Stretching, and enter a percentage in the Time Stretch text box.
- 3 Choose one of the following time stretching options from the pop-up menu, set related options, and then click OK:

Time-Scale Stretch Stretches the clip without affecting pitch. This method is most commonly used for melodic instruments, like piano, bass, and guitar. Because this method bases the stretch on the actual length and duration of the file, use it only to stretch audio that doesn't have well-defined beats, like a synth pad or sustained string section.

Resample (Affects Pitch) Speeds or slows the playback of a clip to fit the new length without maintaining pitch. This setting is commonly used in R&B and hip hop to achieve exaggerated stretching and compressing of drum tracks, creating a lo-fi sound. This setting also works well for vocals, allowing subtle to radical changes in timbre.

Beat Splice Stretches the clip based on beats detected within the file. This setting works only on clips that have very sharp, transient sounds, like drums. If the waveform already has beat markers, select Use File's Beat Markers to use them. Otherwise, select Auto-Find and adjust the default values as needed.

Hybrid Uses the current Time-Scale Stretch settings when you shorten the clip, and uses the current Beat Splice settings when you lengthen it.

To disable time stretching:

- 1 Right-click the time-stretched clip, and choose Clip Time Stretch Properties.
- 2 Deselect Enable Time Stretching.

Inserting empty audio clips

You can insert empty audio clips as placeholders for audio you plan to record later. This technique is particularly helpful when combined with the Punch In command. (See [“Recording audio in Multitrack View” on page 66.](#))

To insert an empty audio clip:

- 1 Select a range in the track display.
- 2 Choose Insert > Empty Audio Clip, and then choose one of the following:
 - In Current Track (stereo)
 - In Current Track (mono)
 - In All Record-Armed Tracks

Revealing hidden clips

If tracks contain overlapping clips, you can reveal hidden clips throughout a session.

To reveal hidden clips:

Choose Edit > Check for Hidden Clips.

Removing and destroying clips

You can remove selected clips from a session and keep their source files available in the Insert menu and in Edit View. Alternatively, you can destroy selected clips to remove them from a session and close their source files.

To remove selected clips:

Choose Edit > Remove Clips.

To destroy selected clips:

Choose Edit > Destroy Clips.

Working with audio tracks

You can record and mix up to 128 tracks in Adobe Audition, and each track can contain as many clips as you need—the only limit is hard disk space. The track controls appear to the left of the track display, and you can resize these controls to be as wide or narrow as you wish. On the Vol, EQ, and Bus tabs, you can access different sets of controls for volume, equalization, and bus properties. Though the wide variety of track controls may seem intimidating at first, the controls for each track are identical, so if you learn one, you've learned them all.



For information about recording on audio tracks, see [“Recording audio in Multitrack View” on page 66](#).



Track controls on the Vol, EQ, and Bus tabs

Using the Track Properties window

In the Track Properties window, you can adjust several settings for the selected track, including volume, pan, output device, and bit depth. Though you can quickly access most of these options in the track controls, the Track Properties window offers channel and bit depth menus, and visual sliders for volume and pan.

To use the Track Properties window:

- 1 Select the track, and then choose Window > Track Properties.
- 2 Set the desired options.

For more information, see [“Track Properties options” on page 197](#).

Track Properties options

The Track Properties window provides the following options:

Track tabs Displays the properties for each track. Click a track in the track display to add a tab for it to the Track Properties window. Click on a tab to adjust its associated track's settings.

Note: If the track tabs aren't visible, click the Restore button (the downward pointing triangle) on the top right of the tab.

Max View button (Upward pointing triangle) Switches to Max View. Click this button to hide the track tabs and expand the controls for the selected track.

Volume slider Adjusts the track's volume, measured in decibels. Drag the slider up (or click the triangle above it) to increase the track's volume; move it down (or click the triangle below it) to reduce the volume. Alternatively, enter a value in the text box above the slider.

Panning slider Adjusts the track's pan, measured in percentage. Drag the slider to the left (or click the triangle at the left of the slider) to reduce the volume of the right channel; drag the slider to the right (or click the triangle to the right of the slider) to reduce the volume of the left channel. Alternatively, enter a value in the text box beside the slider. Enter negative numbers for a left pan and positive numbers for a right pan.

Title Specifies the track name, which corresponds to the name of the track tab.

Output Shows where the track will be output to, such as a sound card or bus.

Record Shows the device (such as a sound card) that the track will record from.

Channel Determines whether the channel is left-only, right-only, or stereo.

Bit Rate Determines whether the track is 16-bit or 32-bit. An Adobe Audition session can contain a mix of 16-bit and 32-bit tracks.

Effects Displays the Effects Rack for the track.

Lock Locks and unlocks the track if it has one or more real-time effects applied to it.

Bus Displays the wet and dry percentages for a track that's been added to a bus. Drag across these text boxes to adjust these percentages.

EQ Displays the low, middle, and high equalization values of the track. Drag across these text boxes to adjust the EQ levels. For more equalization control, choose Window > Track EQ.

Record Arms the track for recording. Click the Record button in the transport controls to start recording.

Solo Solos the track so that only it will be heard on playback. All other tracks are muted.



To solo multiple tracks, hold down Ctrl and press their Solo buttons.

Mute Mutes the output of the track. More than one track in a session can be muted at one time.

Setting track name, volume, and pan

In the track controls, you can name tracks to identify their contents (for example, "Drums"). You can also specify volume and pan settings.



To change track volume and pan over time, use track envelopes. (See ["Automating mixes with clip envelopes" on page 207.](#))

To name a track:

In the track controls, type in the name text box.



Name text box in the track controls

To change track volume or pan:

In the Volume (V) or Pan text box of the track controls, drag to scroll through values.



To change these settings with a slider, right-click the Volume or Pan text box.

Soloing and muting tracks

You can solo tracks to hear them separately from the rest of a mix. Conversely, you can mute tracks to silence them in a mix.

To solo a track:

In the track controls, click the Solo button .



To solo multiple tracks, hold down Ctrl and press their Solo buttons.

To mute a track:

In the track controls, click the Mute button .

Specifying track input and output devices

Using the In and Out buttons in the track controls, you can specify input and output devices for each track. The text on these buttons changes to reflect the device you specify.

When you specify an output device, you can specify either a hardware output or a bus output. Bus outputs let you create submixes of selected tracks (for example, drum tracks), which you can then route to a hardware output. (See [“Using the Bus Mixer” on page 211.](#))



On the EQ tab of the track controls, the In and Out buttons are hidden by default. To reveal these buttons, increase the width of the track controls by dragging the right border.

To specify an input device for a track:

- 1 In the track controls, click the In button.
- 2 From the Device Type menu, choose the device type.
- 3 From the list box, select the input device.
- 4 In the Input Options section, specify the channel and bit depth. (To apply these input options to all tracks, select Same For All Tracks.)

To specify an output device for a track:

- 1 In the track controls, click the Out button.
- 2 Select a hardware output from the Devices list or a bus output from the Busses list.



The list of devices is determined by the devices you designate in the Device Order dialog box. (See [“Designating which devices you want to use” on page 37.](#))

Setting track channel and bit depth

To set the channel and bit depth for a track, use the Track Properties dialog box.



Track Properties dialog box with channel menu revealed and bit-depth menu highlighted

To set channel and bit depth for a track:

- 1 Select the track, and then choose Window > Track Properties.
- 2 Choose an option from the channel and bit depth menus.

Equalizing tracks

You can equalize audio tracks by using either the track controls or the Track Equalizers window. The track controls provide quick access to commonly used equalization settings; the Track Equalizers window provides access to more precise and sophisticated controls.

In the EQ tab of the track controls, the track equalization text boxes show the current low-, mid-, and high-frequency equalization. You can drag across these text boxes to change equalization settings. To switch between two banks of equalization settings, you can use the Eq/A or Eq/B button. For example, you can adjust settings for the Eq/A bank, and then click the button to access the default, unequalized settings for the Eq/B bank. However, if you change settings while Eq/B is active, those settings are preserved. This functionality lets you compare any two settings.

In the Track Equalizers window, you can specify the center frequency and Q range for the low, middle, and high bands. Then, you can use a graph to visually adjust equalization settings.

Note: In the track controls, track equalization text boxes appear in the EQ tab by default. To reveal these fields in other tabs, increase the width of the track controls.



Switching between Eq/A and Eq/B settings

To equalize a track by using the track controls:

- 1 On the EQ tab, drag across the Low (L), Middle (M), or High (H) text box.
- 2 To switch to a different bank of equalization settings, click the Eq/A or Eq/B button. (Double-click the button to copy the current settings to the other bank.)

To equalize a track by using the Track Equalizers window:

- 1 Select the track, and then choose Window > Track EQ. Alternatively, click the EQ tab in the track controls, and then right-click the H, M, or L box.
- 2 Set the desired options.

For more information, see [“Track Equalizers options” on page 201](#).

Track Equalizers options

The Track Equalizers window provides the following options:

Track tabs Display the equalization for each track. When you first select a track in the track display, a tab for that track is added to the Track Equalizers window. Click a tab to adjust the associated track’s settings.

Note: You won’t be able to see the Track tabs if the Max View window display option is selected.

Presets button (P button in upper right corner.) Accesses the EQ Presets dialog box, which lets you save favorite equalization settings, as well as delete EQ presets you no longer want. To add the current EQ settings as a preset, click the Presets button, and then click Add New. Enter a preset name, and click OK. To select a preset, click the Preset button, and click the preset name. To delete a preset, click the Presets button, click Delete, and then click the preset name.



To access EQ presets from the EQ tab of the track controls, right-click the EQ/A or EQ/B button.

Max View/Restore button (Upward or downward pointing triangle.) Switches between max view, which hides all tabs and expands the controls for the selected track, and tabs view, which shows track tabs.

Expand/Collapse buttons (Right or left pointing triangle.) Shows or hides the equalization frequency, amplitude, and Q sliders and text boxes.

Graph Shows the track’s current equalization curve. The x-axis represents frequency, while the y-axis represents amplitude. You can drag the three control points around the graph to adjust the EQ curve. As you do, the sliders move to reflect your changes. The left-most control point initially represents the EQ curve’s low frequency, while the middle and right-most control points correspond to the middle and high frequencies, respectively.

Q boxes Display the low-, mid- and high-frequency Q values. (Q defines frequency bandwidth.) Drag across the boxes to change a Q value.

Band buttons Activate Q boxes for the Low and Hi settings, creating a parametric equalizer.


Horizontal frequency sliders Determine the low, mid, and high frequencies. Move the sliders to the right to increase frequency; move them to the left to decrease frequency. Alternatively, enter values in the text boxes.

Vertical amplitude sliders Boost or attenuate the low, mid, and high frequencies. Move the sliders up to increase amplitude; move them down to decrease amplitude. Alternatively, enter values in the text boxes.

Working with ReWire tracks

To work with ReWire tracks, you must first set up ReWire connections, assigning ReWire outputs to one or more tracks in a session. (See [“Setting up ReWire connections” on page 42.](#)) ReWire tracks offer similar controls to audio tracks. For example, you can quickly change volume, pan, and equalization settings, or you can apply real-time effects. Similarly, you can assign a different device to a ReWire track at any time. Note, however, that saved sessions store only changes made in Adobe Audition; be sure to also save any changes made in the ReWire slave application.

When you synchronize via ReWire, you link the transport controls and timeline of Adobe Audition and the ReWire slave application. For example, if you click the play button in the slave application, Adobe Audition plays the linked session, sending the audio through the outputs specified in the Device Properties dialog box. (See [“Setting properties for audio output devices” on page 38.](#)) You can also, however, preview individual modules in the slave application to hear them independently of the Adobe Audition session. When you do, the modules send audio through the sound card specified in the Sounds and Audio Devices control panel.

 If you notice a timeline offset between Adobe Audition and the slave application, lower the Playback Buffer Size on the Multitrack tab of the Settings dialog box (choose Options > Settings). The default value is 1, but you can enter values as low as 0.1. Because extremely low buffer sizes may cause audio to drop out, you may need to try different values to find one that is acceptable.

To assign a different device to a ReWire track:

- 1 In the track controls, click the RW button.
- 2 For Device Type, select ReWire.
- 3 Select the device, and click OK.

To convert a ReWire track to an audio track:

Right-click the ReWire track, and choose Bounce.

Working with MIDI tracks

Audition can import MIDI files as clips on MIDI tracks. MIDI tracks contain a subset of the controls available for audio tracks: a name text box and controls for solo, mute, and volume. MIDI tracks also, however, contain one unique control: a Map button for assigning MIDI output devices.

Adobe Audition doesn't include MIDI clips in exported mixdown files. You can, however, convert MIDI clips to audio clips by recording the output of a MIDI sound module on an audio track.

Note: In general MIDI terminology, MIDI tracks are instrument tracks in MIDI files. In Adobe Audition, however, MIDI tracks contain MIDI clips, which in turn contain instrument tracks.

Importing and mapping MIDI files

You can quickly import MIDI files into a session and map the resulting clips to specific MIDI devices and channels.



A MIDI track containing a MIDI clip

To insert a MIDI file into a session:

- 1 Right-click either an empty track or a MIDI track, and choose Insert > MIDI.
- 2 Select the MIDI file, and click Open.

To map a MIDI track to a MIDI device and channel:

- 1 In the track controls area, click the Map button.
- 2 From the track list, select the instrument track.
- 3 From the Device pop-up menu, choose the device you want to play the track. (Choose No Output if you don't want a device to play the track.)

The list of devices is determined by the devices you designate in the Device Order dialog box. (See [“Designating which devices you want to use” on page 37.](#))

- 4 From the Channel pop-up menu, choose the MIDI channel you want to send the track to. (Choose No Mapping if you don't want to send the track to a MIDI channel.)

Changing volume, pitch, and tempo for MIDI clips

To musically blend MIDI clips with other clips in a session, you can change master volume, transpose pitch, and shift tempo.

To change the master volume for a MIDI clip:

- 1 Right-click the clip, and choose Set Controller 7.
- 2 Select Set Controller 7, and enter the desired volume in the “Default controller 7 value” text box.

To transpose a MIDI clip:

Choose Edit > Transpose, and then set the desired number of steps up or down. Values range from -12 (one octave down) to +12 (one octave up).

To set the tempo of a MIDI clip:

- 1 Right-click the clip, and choose Set Tempo.
- 2 From the pop-up menu, choose one of the following:
 - Authored to use the tempo stored with the MIDI file.
 - Match Session to match the tempo of the multitrack session.
 - Custom to enter a new tempo in the Tempo text box.

Zooming and playing MIDI clips

You can zoom in or out on a MIDI clip to show more or less MIDI data for individual instrument tracks. If a particular instrument track is best suited to a session, you can instruct Adobe Audition to play only that track. (By default, Adobe Audition plays all instrument tracks in a MIDI clip.)

If you play a session and discover that MIDI clips are incorrectly configured, you can abruptly stop playback of all MIDI tracks while other tracks continue.

To zoom in or out on a MIDI clip:

Right-click the clip, and choose either Zoom In Vertically or Zoom Out Vertically.

To play a specific instrument track in a MIDI clip:

Right-click the clip, and choose Active Track [instrument track name].

To stop playback of all MIDI tracks:

Choose Options > MIDI Panic Button.

Using real-time effects

In Multitrack View, you can apply real-time effects to audio and ReWire tracks. With these flexible effects, you can adjust effects settings as a mix plays. Because real-time effects are nondestructive, you can remove them from a track at any time. You can also change the order of effects to produce a different sonic texture. (For example, you can place Reverb prior to Sweeping Phaser, or vice versa.)



To change an effects mix over time, use clip envelopes. (See [“Automating mixes with clip envelopes” on page 207.](#))

Applying and removing real-time effects

You can apply real-time effects by using either the Organizer window or the Effects Rack dialog box. To remove or reorder these effects, however, you must use the Effects Rack dialog box. You can also use this dialog box to save groups of real-time effects as a preset, which you can quickly apply to multiple tracks.



Clicking the FX button to access the Effects Rack dialog box

To apply a real-time effect to a track:

From the Effects tab of the Organizer window, drag the effect to the track.



If the Organizer window is closed, right-click the FX button in the track controls, choose Rack Setup, and add the effect in the Effects Rack dialog box.

To change settings for a previously applied effect:

- 1 In the track controls, right-click FX, and choose FX Settings.
- 2 Click the tab for the effect, and change settings as desired.

To remove or reorder a real-time effect:

- 1 In the track controls, right-click FX, and choose Rack Setup.
- 2 Select the effect in the Current Effects Rack list, and then do one of the following:
 - To remove the effect, click Remove.
 - To reorder the effect, click either Move Up or Move Down.

To create or apply an effects group preset:

- 1 In the track controls, right-click FX, and choose Rack Setup.
- 2 Do one of the following:
 - To create a preset, click New, and type a name for the preset.
 - To apply a preset, choose it from the Preset menu, and click Apply.

Mixing real-time effects

In the FX mixer, you can change the ratio of dry to wet sound, bypass effects, and combine effects as serial or parallel groups. By default, multiple effects are combined in serial groups, in which the signal travels directly from the output of one effect to the input of the next. In parallel groups, each effect independently receives the dry signal, and the effect outputs are mixed at equal levels.

When you click Serial or Parallel in the FX mixer, mix settings automatically change to achieve the results above. For serial groups, effect inputs are set to 0% of the dry source (specified in the Src text box) and 100% of the previous effect (specified in the Prv text box). Likewise, all effect output sliders are set to zero except for the final slider, which is set to 100%. For parallel groups, effect inputs are set to 100% of the dry source and 0% of the previous effect, while effect output sliders are set to an equal level (33% each for three effects, 25% each for four effects, and so on).

Note: The first effect in the FX mixer lacks Src and Prv text boxes because no previous effect exists.



FX mixer

To mix real-time effects:

- 1 In the track controls, right-click the FX button, and choose FX Mixer.
- 2 Do any of the following:
 - To change the ratio of dry to wet sound that the track outputs, move the Dry Out slider and the effects sliders.
 - To bypass an effect, click Bypass.
 - To combine effects in serial or parallel groups, click either Serial or Parallel.
 - To change the ratio of dry to wet audio that an effect receives, enter percentages in the Src and Prv text boxes. (Src represents the dry sound; Prv represents the output of the previous effect.)



To bypass all real-time effects for a track, right-click the FX button in the track controls, and choose Bypass.

Locking tracks with real-time effects

After you apply real-time effects to a track and edit them, you can lock the track to save processing power for other mixing tasks—an important consideration for complex mixes. Adobe Audition stores locked tracks in the background mix, removing them from the CPU load.

If a track is locked, you can't edit effects, clips, or envelopes it contains. You can quickly unlock the track, however, if you need to change it. Though locking tracks takes a small amount of processing time, unlocking tracks is instantaneous.

To lock or unlock a track that has real-time effects:

In the track controls, click Lock.


Automating mixes with clip envelopes

With clip envelopes, you can automate volume, pan, and effects settings over time. For example, you can automatically increase clip volume during a critical musical passage and later reduce the volume in a gradual fade out. For tracks with real-time effects, you can also automatically change the ratio of dry to wet sound.

Envelopes operate nondestructively, so they don't change the original audio file in any way. If you open an original file in Edit View, for example, you won't hear the effect of any clip envelopes. Envelopes also operate in real-time, so you can edit them as a mix plays.

You can identify envelopes by color and initial position. For example, volume envelopes are green lines initially placed across the top of clips. Pan envelopes are blue lines placed in the center of clips. You edit envelopes by dragging control points on these lines. With volume envelopes, for example, the top of a clip represents 100% of track volume, while the bottom of a clip represents full attenuation (silence). With pan envelopes, the top of a clip represents full left, while the bottom represents full right. If an envelope is too high or low, preventing you from raising or lowering control points, you can rescale it.






Note: Wet/dry mix envelopes have the same initial position as volume envelopes, so you may need to hide one to reveal the other.



Two envelopes in the track display
A. Volume envelope **B.** Pan envelope

To show or hide envelopes:


In the toolbar, click any of the following buttons:

- Show Volume Envelopes .
- Show Pan Envelopes .
- Show Wet/Dry Mix Envelopes .
- Show FX Parameter Envelopes .
- Show Tempo Envelopes .



You cannot edit tempo envelopes, which display the tempo of MIDI clips. For information about changing MIDI tempo, see [“Changing volume, pitch, and tempo for MIDI clips” on page 204.](#)

To edit a clip envelope:

- 1 In the toolbar, click the Edit Envelopes button .
- 2 Select the clip, and then do any of the following:
 - To add a control point, click the envelope.
 - To remove a control point, drag it off the clip.
 - To move a control point, drag it. (To maintain time position, hold down Shift while dragging.)
 - To move all control points up or down by the same percentage, hold down Ctrl while you drag.
 - To move all control points up or down by the same amount, hold down Alt while you drag. (This option retains envelope shape, restricting movement to the limits defined by the highest and lowest control points.)

Note: With MIDI clips, volume envelopes control MIDI velocity, which usually represents the force with which a note is struck. However, some synthesizers are programmed so that velocity changes pitch or harmonic content.

To clear all control points for an envelope:

Right-click the clip containing the envelope, and choose Envelopes > [envelope type] > Clear Selected Points.

To use spline curves for an envelope:

Right-click the clip containing the envelope, and choose Envelopes > [envelope type] > Use Splines.

To rescale a volume envelope:

- 1 Right-click the clip containing the envelope, and choose Rescale Volume Envelopes.
- 2 Enter the number of decibels by which you want to raise or lower the envelope. Possible values range from -40 to 40. Negative values raise envelopes and lower clip volume by an equal amount; positive values do the opposite.



You can also rescale all volume envelopes in a session. See [“Mixing options for sessions” on page 185.](#)

Using the Mixers window

The Mixers window consists of the Track Mixer and Bus Mixer tabs, as well as a slider that controls the master volume of your session.

The Track Mixer tab mimics a real-world mixing console. It gives you an alternative view of a session, providing a broader overview than the track display, especially if you're working with more than a handful of tracks at once.

The Bus Mixer tab lets you create, configure, and control up to 26 buses. With buses, you can group related tracks and collectively adjust volume or apply real-time effects.



Even if you don't use the Mixers window, consider docking and resizing it so that only the master volume slider is visible. You can use this slider to quickly optimize the overall volume of a mix.

Using the Track Mixer

The Track Mixer provides another method for viewing the tracks in a session. Though it lacks the waveforms, clips, and envelopes visible in the track display, the Track Mixer lets you view and edit more tracks simultaneously.



To automate volume and pan changes over time, use clip envelopes. (See [“Automating mixes with clip envelopes” on page 207.](#))



The Track Mixer

A. Track controls **B.** Control display buttons **C.** Scroll bar

To use the Track Mixer:

- 1 In Multitrack View, choose Window > Mixer.
- 2 Click the Track Mixer tab, and set the following options:

Control display buttons Lets you customize the appearance of the Tracks Mixer. Each of the six buttons—Out, Bus, FX, EQ, Pan, and M/S—displays a different track control.

- Out shows and hides the Out buttons.
- Bus shows and hides the Wet and Dry text boxes for buses.
- FX shows and hides the FX and Lock buttons.
- EQ shows and hides the three equalization text boxes (H, M, L).
- Pan shows and hides the Pan controls.
- M/S: Shows and hides the Mute and Solo buttons.

Out 1 Opens the Adobe Audition Playback Devices window, which lets you assign the output properties for the selected track. The button's label changes to reflect the output device (for example, Device 2 or Bus C).

Wet and Dry text boxes Control the ratio of processed to original signal that tracks output to buses. To change these values, either enter a value or drag across the text boxes. (You can also change these values on the Bus tab of the track controls.)



To output a track to a bus, see [“Specifying track input and output devices” on page 199](#).

FX Opens either the Effects Rack (if the track doesn't yet have an effect assigned to it, regardless of what effects may be assigned if the track is part of a bus) or the dialog box for whatever effect is assigned to the track.

Lock Locks or unlocks a track. If the Lock button is disabled, the track has no effects.

H, M, L Show the amplitude of high, mid, and low equalization frequencies applied to the track. To change one of these values, drag across the text box. Dragging to the right increases the value, while dragging to the left reduces it.

Pan controls Adjust the balance of each track. To use the control that looks like an asterisk, drag it to one of three positions: hard left, zero pan, and hard right. The Pan text box provides a more precise way of adjusting pan. To change a pan value, drag across the text box to the left or right.

Mute and Solo buttons Let you mute or solo a track. Click the Mute button for as many tracks as you like to turn off their output. Click the Solo button to solo the track. To solo multiple tracks, hold down the Ctrl key as you click the Solo buttons.

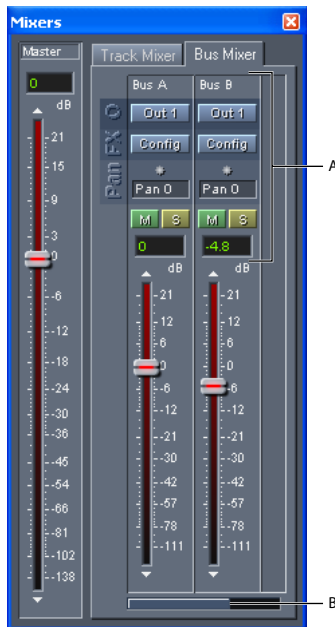
Track faders Adjust the track's relative volume in the mix. Move the slider up (or click the triangle above it) to increase the volume; move it down (or click the triangle below it) to reduce the volume. Alternatively, enter a value (in decibels) in the text box above the slider.

Scroll bar Lets you scroll from tracks 1 to 128 and any point in between.

Using the Bus Mixer

Adobe Audition gives you the ability to organize multiple tracks into buses, which are especially useful for grouping related tracks and collectively adding real-time effects or adjusting volume. For example, you can output four tracks of background vocals to one bus, and then apply one reverb effect to that bus. (Individually applying the same reverb to each vocal track would inefficiently drain CPU resources.) You can create up to 26 buses.

After you create and configure a bus, you can output tracks to it. (See [“Specifying track input and output devices” on page 199](#).) Then, on the Bus tab of the track controls, you can change the ratio of wet to dry sound that tracks send to assigned buses.



The Bus Mixer

A. Bus controls **B.** Scroll bar

To create and configure a new bus:

- 1 In Multitrack View, choose Window > Mixer.
- 2 Click the Bus Mixer tab, and then click New in the right-most Bus channel.
- 3 In the Bus Properties dialog box, enter a name in the Friendly Name field, and select an output device.
- 4 In Installed Real-Time Effects list, select effects for the bus, and click Add.
- 5 Click OK.

To mix and reconfigure buses:

- 1 In Multitrack View, choose Window > Mixer.
- 2 Click the Bus Mixer tab, and then set the following options:

Out Opens the Bus Properties dialog box, where you can specify a different output device or combination of effects.

Config Opens the configuration window for the selected bus. Here you can access the parameters for each effect added to the bus. Sliders also let you adjust the volume of all effects in the bus, as well as the desired Dry Out level.

- Click Serial to connect the effects on the bus in sequence, with the output of one effect connected to the input of the next.
- Click Parallel to connect the effects on the bus separately, mixing only their outputs together.
- Click Rack Setup to open the Properties dialog box for the bus.

Pan controls Adjust the balance of each track. To use the control that looks like an asterisk, drag it to one of three positions: hard left, zero pan, and hard right. The Pan text box provides a more precise way of adjusting pan. To change a pan value, drag across the text box to the left or right.

Mute and Solo buttons Let you mute or solo a bus. Click the Mute button for as many buses as you like to turn off their output. Click the Solo button to solo the bus. To solo multiple buses, hold down the Ctrl key as you click the Solo buttons.

Bus faders Adjusts the relative volume of the bus in the mix. Move the slider up (or click the triangle above it) to increase the volume; move it down (or click the triangle below it) to reduce volume. Alternatively, enter a value in decibels in the text box above the slider.

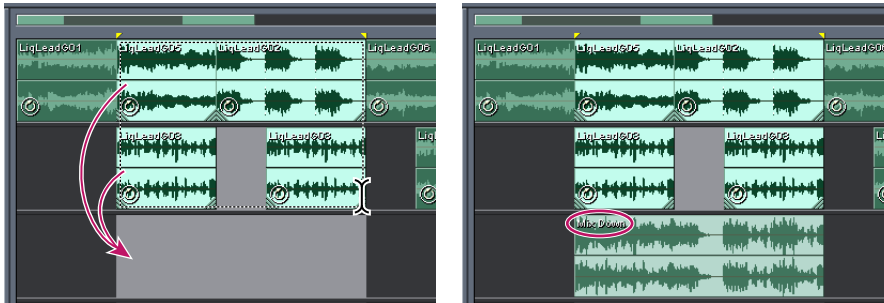
Scroll bar Lets you scroll through buses if all of them don't fit in the Bus Mixer tab.

Mixing down ReWire tracks and specific audio clips

You can mix down ReWire tracks and specific audio clips to a new file that opens in Edit View, an empty track in the current session, or a track in a CD project.



You can also mix down entire sessions, exporting them in a variety of formats. (See [“Saving and exporting sessions” on page 238.](#))



Mixing down specific audio clips to an empty track

To mix down ReWire tracks and specific audio clips:

- 1 Select any audio clips you want to mix down.
- 2 Select the range you want to mix down.



If you want to mix down complete clips, you can skip step 2 if the session doesn't contain ReWire tracks.

3 From the Edit menu, choose Mix Down To File, Mix Down To Empty Track, or Mix Down To CD Project, and then choose one of the following:

- All Audio Clips to mix down ReWire tracks and all audio clips.
- Selected Audio Clips to mix down ReWire tracks and selected audio clips.
- All Audio Clips (Mono) to mix down ReWire tracks and all audio clips in mono.
- Selected Audio Clips (Mono) to mix down ReWire tracks and selected audio clips in mono.

Using Loops

About loops

Loop-based song creation has recently cropped up in nearly all musical circles. From best-selling pop, rap, and hip hop songs to the alternative, adult contemporary and jazz realms, using loops, even as basic rhythm tracks, is a very appealing and modern technique for making music. With Adobe Audition, you can create your own loops or access any of the thousands supplied in the Adobe Audition Loop Library.

Loops typically contain one to two bars of music. Most pop and rock music follows a 4/4 time signature, meaning that one bar has four beats, two bars have eight beats, and so on.

With loops in Adobe Audition, you can do the following:

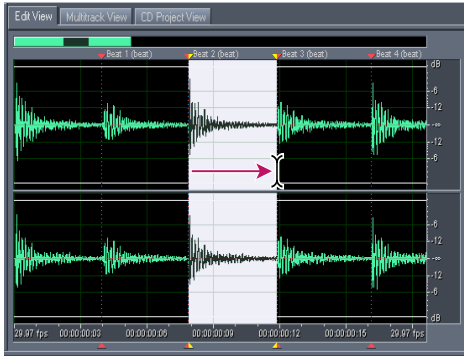
- Change the pitch and timing of loops independently of each other, so you can easily incorporate the same loop into many different Adobe Audition sessions and musical compositions.
- Quickly and easily add or subtract repetitions of a loop by dragging with the mouse. (With snapping enabled, this method applies even to individual beats within a loop. For example, you can drag to create 1.5 repetitions and end precisely on a snare hit at a loop's midpoint.)
- Snap other audio clips to loop end-points and beats within the loop.

Working with loops is typically a three-step process, in which you select part of a waveform, specify its properties in Edit View, and then use the resulting loop in compositions in Multitrack View.



Defining loops

To come up with a good loop, you first need to select and save a waveform that, when played over and over, repeats precisely on a beat. This process is called *defining* a loop. Although defining on a beat isn't absolutely required, doing so makes loops more useful because you can combine them in rhythm with other loops.




Selecting a waveform that starts and ends on a clear beat helps make for a good loop.

To define a loop:

- 1 Open the waveform from which you want to define the loop.
- 2 Switch to Edit View.
- 3 Choose Edit > Auto-Cue > Find Beats And Mark. Set up the dialog box to find the beats in the waveform, and click OK. See [“Setting cues automatically” on page 85](#) for more information on the Find Beats And Mark command, and see [“Finding beats” on page 78](#) for information on setting up Adobe Audition to find beats.

The beats in the waveform are now indicated, helping you select a start and end that lands on the beat.

- 4 Choose Edit > Snapping > Snap To Zero Crossings. (See [“Snapping” on page 78](#).) This step makes your selection snap to places in the waveform that have zero amplitude, preventing audible noise at the beginning and end of the loop.
- 5 Select the part of the waveform you want to define, typically starting and ending on a beat.
- 6 Click the Play Looped button  to repeatedly play your selection.
- 7 Adjust the start and end of the selection until just the material you want is selected.
- 8 Choose View > Display Time Format > Edit Tempo to specify detailed tempo information for the loop. (See [“Calculating the tempo of selected ranges” on page 216](#).)
- 9 Choose Edit > Copy To New. This step copies the selected area to a new file where you can set the loop's properties. (See [“Setting permanent loop properties in Edit View” on page 216](#).)

Calculating the tempo of selected ranges

In both Edit View and Multitrack View, you can calculate the tempo of a selected range by using the Edit Tempo command. This command lets you quickly determine loop tempo in Edit View or change session tempo in Multitrack View. It also lets you change the beats-per-minute (bpm) value for horizontal rulers in the Bars and Beats time format.

To calculate the tempo of a selected range:

- 1 Choose View > Display Time Format > Edit Tempo.
- 2 Set the following options, and click OK:

Beats Highlighted/Bars Highlighted Specifies the number of beats or bars highlighted in the selection according to the Bars and Beats format. This number will probably be wrong initially, because you haven't defined the tempo yet. In this case, enter the correct number of bars to use for extracting tempo information.

Extract Calculates tempo information from the highlighted selection, and fills in the Beats per Minute and Offset values. Before clicking Extract, make sure to enter a value for Beats per Bar.

Current Beat At Defines the bar and beat information for the selection's starting point (or the current cursor position if no selection has been made). Adobe Audition assumes that this represents a downbeat. Changing this value updates the Song Start value based on the current tempo settings.

Reset 1:1 To Cursor Changes the Current Beat At value to 1:1.00.

Song Start Represents the number of milliseconds before the measure 1:1.00 begins. This value is for information only.

Beats Per Minute Displays the number of beats that occur in a one-minute interval. You can calculate this value by clicking Extract.

Beats Per Bar Assigns the number of beats that make up one measure/bar. For instance, enter **4** for 4/4 time, **6** for 6/8 time, and the like.

Beat Length Specifies the value of the beat. For instance, enter **2** for a half note, **4** for a quarter note, and **8** for a sixteenth note.

Ticks Per Beat Specifies the number of sections each beat is divided into, or the value after the decimal point. You can enter a number between 2 and 3600. For instance, if you enter 32 ticks per beat, then a time setting of 4:2:16 represents an eighth note (a note halfway) between beats 2 and 3 in 4/4 time.

Setting permanent loop properties in Edit View

After you define a loop, you can set permanent loop properties so it works well with other clips in a session. The Loop Info tab of the Wave Properties dialog box lets you specify these properties, such as number of beats, default tempo, and musical key. Setting permanent loop properties makes a loop far easier to work with in Multitrack View.

Loop properties that you set in Edit View are saved with the file and are permanent. Loop properties that you set in Multitrack View are saved with the session and aren't permanent. In addition, session-based loop properties in Multitrack View override Edit View loop properties. See [“Setting impermanent loop properties in Multitrack View” on page 218](#) for more information.

To set loop properties in Edit View:

- 1 Choose View > Wave Properties.
- 2 Click the Loop Info tab.
- 3 Set any of the following options, click OK, and then save the file:

Loop Tells Adobe Audition that the file is a loop. If the file is inserted into a Multitrack session, looping is enabled automatically for that audio clip.

One Shot Indicates that the file plays once rather than repeats like a loop.

Number Of Beats Specifies the number of beats in the loop. Adobe Audition attempts to detect and specify the number of beats for you, but you can adjust the value if necessary.

Tempo Specifies the number of beats per minute in the loop. Adobe Audition calculates this value automatically based on Number Of Beats. Don't worry if the value isn't a whole number—for example, 80.4 instead of 80—after you loop the file, Adobe Audition can stretch it to whatever tempo you want.

Key Specifies the loop's key, so that if you create a session and want to adjust the key of all audio clips globally, Adobe Audition has a reference for each file. If a loop file is a drum track, choose Non-Voiced. This option is especially important if you plan to change the key of multiple loops in a session, because you won't want to pitch-shift a drum track to the key of E (for example) if it has no key to begin with.

Find Nearest Scans the loop to locate the nearest key. This option works best with monophonic files (that is, solo instruments). Because many keys share the same notes in the scale, you can think of this setting as root note for transposition.

Stretch Method Specifies how (if at all) the loop stretches to match the session's tempo. Choose one of the following settings:

- Fixed Length (No Stretching) causes the loop to play at its native tempo no matter what the session tempo is set to. If a session has multiple loops of different tempos, and each is set to Fixed Length, no two loops will seamlessly match in tempo. This setting is useful if you plan to insert and loop a file in a session where you don't plan on doing any type of time stretching or pitch shifting. The most common uses for this setting is inserting a pattern over live music or using one to underscore live vocals.
- Time-Scale Stretch stretches the file (just like the Stretch effect) to match the tempo of the session. Corresponding options are Quality (High, Medium, or Low), Frame Size (the number of splices per beat), and the percentage of Frame Overlapping. This method stretches a file based on its actual length, so you should use it if you loop something like a synth pad or a sustained string section (which don't have actual beats, per se). This method is most commonly used for "tonal" instruments, like piano, bass, and guitar.
- Resample (Affects Pitch) resamples the loop to match the session's tempo, affecting the pitch. High, Medium, and Low Quality options are available. This method is commonly used in R&B and hip hop tracks, primarily because you can achieve exaggerated stretching and compressing of files. If loops set to Resample are time stretched, their pitch changes. This setting is most commonly used on drum tracks to create a lo-fi, dirty, phat kind of sound. It can also work well on voice and voiceovers if you're trying to change the sound and timbre of a speaker's voice.

- Beat Splice loops the file based on beats detected in it, similarly to the Find Beats And Mark command. (See [“Defining loops” on page 215](#).) This setting works only on loops that have very sharp and short sounds, like drum tracks. If the waveform already has beat markers, you can select Use File’s Beat Marks to use them. Otherwise, Auto-Find Beats is selected. If necessary, you can change the corresponding default values of 10 dB and 9 milliseconds to find the beat.
- Hybrid uses the current Time-Scale Stretch settings if you lower the bpm (beats per minute), and it uses the current Beat Splice settings if you raise the bpm.

Setting impermanent loop properties in Multitrack View

Loop properties that you set in Multitrack View are saved with the session and aren’t permanent, but they override any permanent loop properties you’ve set in Edit View. (See [“Setting permanent loop properties in Edit View” on page 216](#).)

By default, changes made to a looped audio clip in Multitrack View affect only that clip, unless Adjust All Loop-Enabled Clips That Use This Wave is selected in the Loop Properties dialog box.

To set impermanent loop properties in Multitrack View:

- 1 Select an audio clip.
- 2 Choose Edit > Loop Properties.
- 3 In the Audio Clip Looping dialog box, set the following options, and click OK:

Enable Looping Sets the file so that you can loop the audio clip by dragging its right edge.

Simple Looping (No Gaps) Makes the audio clip loop continuously, with no spaces between looping instances.

Repeat Every X Seconds Repeats the loop at the number of seconds you specify. If loop information is already entered for the audio clip, proper values for Repeat Every X Seconds and Repeat Every X Beats are entered automatically so that the audio clip loops continuously at the proper tempo. If you change the Repeat Every X Seconds value, Adobe Audition ignores the tempo and stretches the file to finish its loop in the specified number of seconds. Normally, you should select this option and enter the number of beats in the Source Waveform Information area.

Repeat Every X Beats Repeats the loop at the number of beats you specify. If loop information is already entered for the audio clip, proper values for Repeat Every X Seconds and Repeat Every X Beats are entered automatically so that the audio clip loops continuously at the proper tempo. If you change the Repeat Every X Beats value, Adobe Audition stretches the file to finish its loop in the specified number of beats. However, you’ll generally want to select Repeat Every X Beats and enter the number of beats in the Source Waveform Information area.

Follow Session Tempo Plays the loop at the session’s tempo instead of its native tempo. For example, if you play a 100 bpm loop in a 120 bpm session, the loop is stretched to 120 bpm. Selecting this option disables the BPM text box. If you don’t select this option, the loop plays at the tempo specified in the BPM text box.

Lock Position To Tempo Locks the left edge of the audio clip to the bar/beat. If you change tempo, the audio clip moves so that it starts at the same beat. Normally, you should select this option if you stretch to tempo. In addition, you can select this option for a one-shot clip that's not a loop (like a thunder clap or a gong) if you want it to start in time with other music that's aligned to the session's tempo.

Source Waveform Information Specifies settings for the source waveform. (See [“Setting permanent loop properties in Edit View” on page 216.](#))


Tempo Matching Specifies settings for matching the tempo of the loop to the rest of the sound file you're working with. (See [“Setting permanent loop properties in Edit View” on page 216.](#))

Transpose Pitch Transposes the pitch of the looped clip by the specified number of half-steps. Positive numbers raise the pitch, and negative numbers lower it.

Adjust All Loop-enabled Clips That Use This Wave Globally changes the settings for all clips that reference the same waveform. For example, if you insert the same loop file into Multitrack View four times, and you then adjust the loop properties on one of its clips, the other three instances of the loop in the session are adjusted, too.

Setting the tempo, time signature, and key for sessions

The Session Properties window lets you specify the tempo, time signature, and key for loops in a session. All loop-enabled clips automatically adjust to match new settings; regular clips are unaffected.

 To preview loop files at the tempo and key of a session, select either the Loop option in the Insert Audio dialog box, or the Follow Session option in the Files tab of the Organizer window. (See [“Inserting audio files into multitrack sessions” on page 58](#) and [“Previewing audio by using the Organizer window” on page 69.](#))

To set the tempo, time signature, and key for a session:

1 In Multitrack View, choose Window > Session Properties if the window isn't visible.

2 Set any of the following options:

Tempo Specifies the tempo of the session, measured in beats per minute.

Beats/Bar Specifies the number of beats per bar.

Key Specifies the session's key.

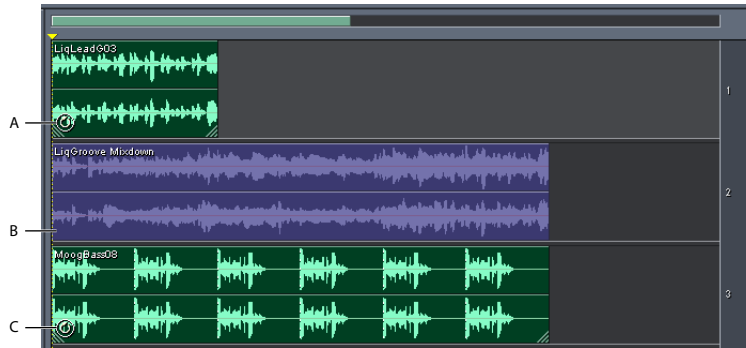
Time Specifies the session's time signature. Choosing a different time signature automatically updates the Beats/Bar setting.

Advanced Opens the Advanced Session Properties dialog box so that you can set advanced properties for a session, such as a time offset, a customized metronome, and notes about the session. (See [“Setting advanced session properties” on page 184.](#))

Metronome Toggles the built-in metronome on and off. (See [“Metronome options for sessions” on page 186.](#))

Working with loops in the track display

After you add loops to a multitrack session, you can edit them in the track display, extending them to repeat as needed, and synchronizing them with the beat of the music.



Loops in the track display

A. Single loop **B.** No loop **C.** Extended (repeated) loops. Even though loop files are short, you can extend them to repeat as many times as needed.

To synchronize loops to musical beats:

1 Choose View > Display Time Format > Bars And Beats to change the format of the ruler to bars: beats:ticks per beat. This format makes it easier to visually line up loops with musical beats. (See [“Monitoring time” on page 63.](#))


2 From the Edit > Snapping submenu, choose any of the following

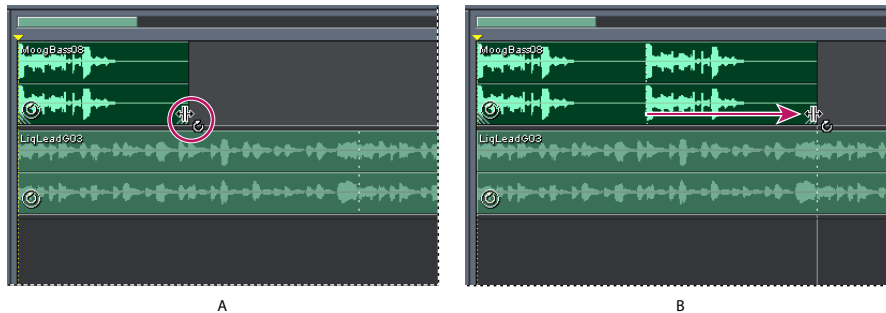
- Snap To Ruler (Coarse) to snap loops to the beats within bars. Use this option if you work with 1/4 or 1/2 bar loop files. (See [“Snapping” on page 78.](#))
- Snap To Clips to snap loops to the start or end of audio clips. (See [“Snapping clips to loop endpoints and other clips” on page 189.](#))
- Snap To Loop Endpoints to snap loops to the start or end of other loops. (See [“Snapping clips to loop endpoints and other clips” on page 189.](#))



Also consider snapping non-loop-enabled audio clips to the beat and each other, so that all clips are aligned. You can snap the current-time indicator, too.

To extend or shorten a loop-enabled clip:

- 1 Select the clip, and then position the pointer over the bottom left or right handle—the loop editing icon  appears.
- 2 Drag the handle to extend the loop the desired number of bars. Depending on how far you drag, you can make the loop repeat fully or partially. For example, you might drag a loop that is one bar long so that it extends 3-1/2 bars, ending on a beat within the loop. As you cross each bar, a white vertical line appears in the clip. This is the snap-to line, indicating perfect alignment to beats in other tracks.



Extending a loop

A. Moving the cursor to activate the loop. **B.** Dragging the loop, with snap to lines indicating how the loop snaps to the beats in other tracks

Working with Video

About working with video

With Adobe Audition, you can improve the sound of any video project. For example, if you need to improve the audio quality of an existing soundtrack, you can use Edit View to quickly restore and enhance the audio. Or, if you want to create elaborate soundtracks with flexible, real-time mixing tools, you can use Multitrack View to preview video, add audio and MIDI tracks, and export entirely new soundtracks. (See [“Importing audio and video from video files” on page 223.](#))

For maximum flexibility, you can combine Adobe Audition with Adobe Premiere Pro and After Effects. Tight integration between these products lets you quickly remix a soundtrack as the needs of a video project change over time. (See [“Working with Adobe Premiere Pro and After Effects” on page 222.](#))

Working with Adobe Premiere Pro and After Effects

If you use Adobe Premiere Pro or After Effects, you can easily remix and edit soundtracks in Adobe Audition. To do so, first configure Adobe Audition to link session files with exported audio mixdowns in WAV format. Once these files are linked, you can select an imported mixdown file in Adobe Premiere Pro or After Effects, and then remix the related session in Multitrack View, or edit the mixdown file in Edit View.

To link session files with exported audio mixdowns in WAV format:

- 1 Choose Options > Settings, and then click the Data tab.
- 2 Select Embed Project Link Data For Edit Original Functionality, and then click OK.
- 3 When you export mixdown files, select Save Extra Non-Audio Information in the Export Audio dialog box.

To remix or edit a mixdown in an Adobe Premiere Pro or After Effects project:

- 1 In the Adobe Premiere Pro or After Effects project, select the mixdown file.
- 2 Choose Edit > Edit Original.
- 3 Select one of the following, and then click OK:
 - Launch The Audition Multitrack Session Which Created This File.
 - Insert This File Into Audition's Edit View.
- 4 Remix the linked session in Multitrack View, or edit the mixdown file in Edit View.
- 5 Overwrite the original file by doing one of the following:
 - In Multitrack View, choose File > Export > Audio, and specify the same name and location as the original file.
 - In Edit View, choose File > Save.



Importing audio and video from video files

In both Edit View and Multitrack View, you can import audio data from a video file in AVI, MPEG, or WMV format. This approach is useful for soundtrack editing that doesn't require a video preview, or for readapting soundtracks for audio-only mediums, such as radio or CD.

Only in Multitrack View, however, can you import both audio and video data from a video file. This approach lets you precisely synchronize audio with a video preview. Note, however, that a multitrack session can contain only one video clip at a time.

To import audio data from a video file:

Do one of the following:

- In Edit View, choose File > Open Audio From Video.
- In Multitrack View, select a track, position the current-time indicator at the desired insert point, and then choose Insert > Audio From Video.

To import audio and video data:

In Multitrack View, select a track, position the current-time indicator at the desired insert point, and then choose Insert > Video.

Working with video clips

When you import a video file into a multitrack session, video data becomes a video clip on the selected track, and audio data becomes an audio clip on the track below. You can select and move video clips like other clips. (See [“Selecting and moving clips” on page 188](#).) Note, however, that you can also move a video clip independently from the audio clip containing the original soundtrack; to keep related video and audio clips synchronized, group them. (See [“Grouping clips” on page 188](#).)

To synchronize audio and video, you can snap other clips and the current-time indicator to individual frames in a video clip. You can also magnify the session display to view more thumbnails in a clip. These thumbnails serve only as a general guide; for frame-accurate synchronization, use snapping.

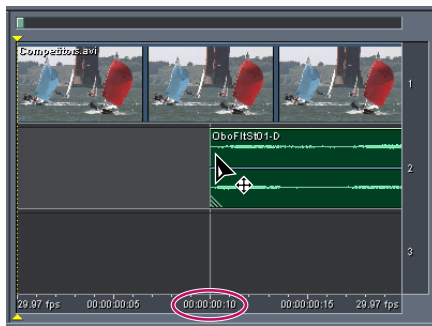
Note: *Thumbnails don't appear for MPEG-2 video clips.*

To snap to individual frames in a video clip:

- 1 Choose View > Display Time Format, and select the SMPTE time format that corresponds to the frame rate of the clip.
- 2 Choose Edit > Snapping > Snap To Frames.

To view more thumbnails in a video clip:

Horizontally magnify the session display. (See [“Zooming”](#) on page 25.)



Snapping to a video frame that falls within a video thumbnail

Previewing video

In the Video window, you can preview video clips as a multitrack session plays to precisely synchronize a soundtrack with specific video events such as scene changes, title sequences, or special effects. You can customize the preview to optimize it for your monitor size and system speed. For example, you can enlarge the preview to fit a resized Video window or lower the preview quality to increase performance.



The floating Video window in Multitrack View

To hide or show the Video window:

In the toolbar, click the Hide/Show Video Window button .

To automatically show the Video window when you insert a video file:

Right-click the Video window, and select Auto Show Video.

To customize the video preview:

Right-click the Video window, and select any of the following:

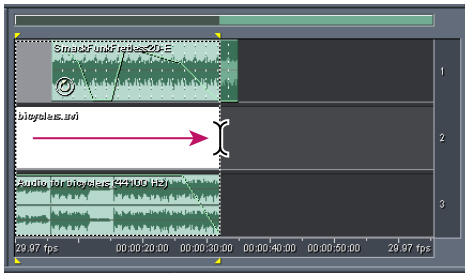
- A zoom percentage to zoom in or out.
- Best Fit to fit the preview to the window.
- Maintain Aspect Ratio to maintain that ratio when you resize the window.

- Integer Factor Sizing to constrain the preview to ratios such as 1/2, 1/1, and 2/1 when you resize the window. This option avoids complex resampling, producing a sharper image and increasing performance.
- Low Quality to lower the preview quality.

Note: Video quality settings take effect when you next import a video clip. To apply a new quality setting to the current clip, close it, and reimport it into the session.

Preparing video mixdowns for export

In Multitrack View, you can export video mixdowns in AVI format. Video mixdowns combine video clips with audio clips that exist in the same area of the timeline, creating a new soundtrack. Prior to exporting a video mixdown, you can preview it to ensure that it will sound as you expect and, if not, edit the session as desired.



Selecting the start and end points of a video mixdown

To preview a video mixdown:

- 1 Choose Edit > Snapping > Snap To Frames.
- 2 In the session display, select an area that extends from the beginning to the end of the video clip.
- 3 Play the session, and then do one of the following:
 - If the soundtrack doesn't sound as you expect, edit the session as desired, and then repeat steps 2 through 3. (For example, if part of an audio clip is omitted, move the entire clip into the selected area.)
 - If the soundtrack sounds as you expect, export a video mixdown. (See ["Exporting mixes to video" on page 239](#).)

💡 You can also use this procedure to export an audio mixdown that you combine with video in a video application, such as Adobe Premiere Pro. Though video mixdowns are limited to stereo audio and AVI format, audio mixdowns support stereo and surround sound in a variety of formats. For more information, see ["Exporting mixes to audio" on page 239](#) and ["About surround sound" on page 226](#).

Creating Surround Sound

About surround sound

With surround sound, heard in many popular movies, you can pan an audio mix around the room. Adobe Audition supports 5.1 surround sound, which requires five speakers, plus one low frequency subwoofer (LFE). To properly preview a 5.1 surround-sound mix, your computer must have a sound card with at least six outputs, and the speakers must be connected and positioned as follows:

- Output 1: Front left speaker.
- Output 2: Front right speaker.
- Output 3: Front center speaker.
- Output 4: LFE.
- Output 5: Left surround speaker.
- Output 6: Right surround speaker.

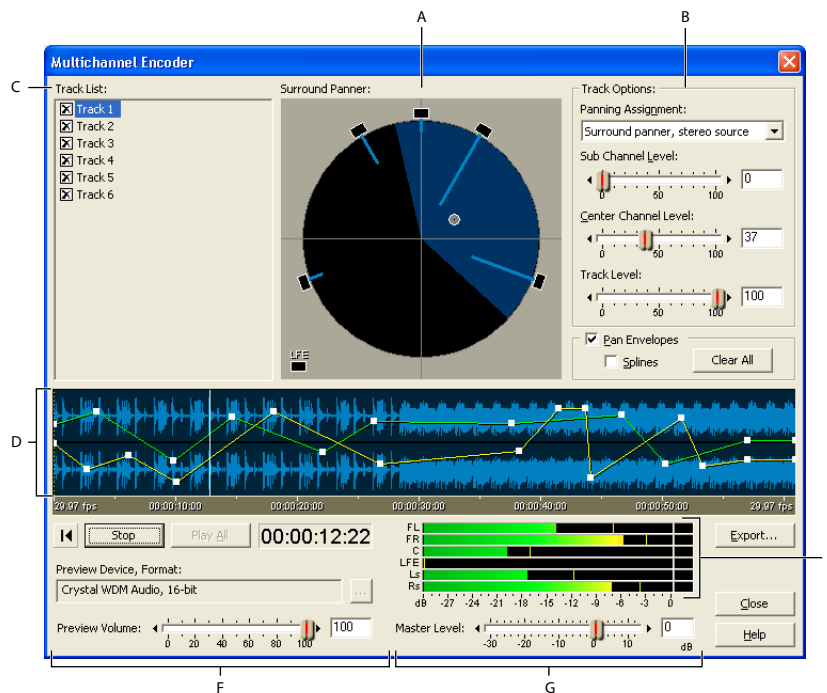
Adobe Audition lets you create and export 5.1 surround sound in a multichannel session by using the Multichannel Encoder dialog box. With this dialog box, you can individually pan each track of a multitrack session to your multichannel setup, preview the current mix, and export the session. You can export your session as six mono WAV files, as one interleaved 6-channel WAV file, or as a Windows Media 9 Pro (WMA) file for use with an external multichannel encoder such as a Dolby or DTS encoder.

Using the Multichannel Encoder

The Multichannel Encoder dialog box contains several options and controls that let you select tracks and bus outputs, precisely pan audio and adjust volume levels, zoom the waveform display, and preview the project.



To achieve proper 5.1 surround-sound preview playback from the Multichannel Encoder, you need a sound card that offers at least 6-channel analog output, a special interleaved driver that's compatible with the Microsoft DirectSound multichannel format, and Microsoft DirectX 8.0 or later. (Direct X 8.0 is installed by default as part of the Adobe Audition installation; updates are available from the Microsoft Web site.) If your system does not meet these requirements you may receive a warning message and your Play Track and Play All buttons will not be accessible.



The Multichannel Encoder dialog box

A. Surround Panner **B.** Track options **C.** Track List **D.** Waveform display with pan envelopes
E. Output meters **F.** Preview controls **G.** Master Volume control

To use the Multichannel Encoder:

- 1 Open an existing Adobe Audition session, or create a new session in the Multitrack window.
- 2 Once all your tracks are added, achieve a basic stereo mix balance with your desired track volume, stereo pan, and FX settings.
- 3 Choose View > Multichannel Encoder.
- 4 In the Track List, select the tracks and bus outputs you want to pan and export. (See [“Selecting tracks and buses in the Multichannel Encoder” on page 228.](#))
- 5 Under Track Options, specify the Panning Assignment and set the controls as desired. (See [“Assigning the panning source” on page 228](#), [“Using the Surround Panner” on page 229](#), and [“Automating the pan envelope” on page 230.](#))
- 6 Set the volume for the tracks. (See [“Adjusting volume levels” on page 231.](#))
- 7 Preview the panned tracks. (See [“Previewing the multichannel project” on page 232.](#))
- 8 Export the session. (See [“Exporting surround-sound files” on page 234.](#))

Panning tracks and buses for surround sound

Using the Multichannel Encoder, you can pan any of the tracks and buses in your session for surround sound. By panning sound between the six surround-sound speakers, you can make the sound appear to come from anywhere around the listener.


Selecting tracks and buses in the Multichannel Encoder

All of the tracks used in your current multitrack session appear in the Track List in the Multichannel Encoder. If you uncheck a track it is removed from the multichannel preview and is not included when you export the multichannel project.

If you have routed a track to a bus, the bus, instead of the track, will appear in the track list. You can select and pan the bus output as one mono or stereo signal. Additionally, you can access the “dry” track signal and pan it separately as well. (See [“Using the Bus Mixer” on page 211](#) for more information on setting up a bus.)

To select a track to pan for surround sound:

In the Multichannel Encoder dialog box, click the name of the track.



The Track List with Track 4 selected

To access a track assigned to a bus in the Track List:

- 1 Close the Multichannel Encoder.
- 2 In the track controls in Multitrack View, click the Bus tab for the track you want to access.
- 3 Increase the Dry value so that it is greater than zero.
- 4 Reopen the Multitrack Encoder.

Assigning the panning source

At the top right of the dialog box is the Panning Assignment list where you can choose to either use the Surround Panner to position your track sound source or make fixed panning assignments for your track.

To specify the panning assignment:

Choose one of the following options from the Panning Assignment list at the top right of the Multichannel Encoder dialog box:

Surround Panner, Stereo Source Uses the Surround Panner to position your sound source. (See [“Using the Surround Panner” on page 229](#).) It also keeps your stereo left and right signals from your track discrete when panning in the sound field. For example, if your track includes a stereo file, the left stereo signal is sent to the Front Left and Left Surround channels, and your track’s right signal is sent to the Front Right and Right Surround channels. The Center channel always receives a summed to mono (L + R) signal. Therefore, as you pan in the five channel sound field, these stereo sources retain their “stereo image” while being routed to the multiple channels.

Surround Panner, Summed To Mono Lets you use the Surround Panner to position your sound source. However, this option always sums the track's signal to a mono signal. In this mode, panning the sound source to any location in the sound field results in the summed mono signal being fed to all channels.

Lfe Only Sends the entire track signal to the LFE (subwoofer) channel. Your monitoring system applies the proper crossover frequency cutoff for reproducing the audio sent to the LFE channel. Typically, most LFE components in 5.1 surround playback systems are set to a cutoff of < 80 Hz or < 120 Hz. The Multichannel Encoder itself does not apply any filter to the LFE channel audio.

FL + FR, Stereo Sends the selected track's signal as a stereo source directly to only the Front Left and Front Right speakers in a 50/50 stereo balance.

Ls + Rs, Stereo Sends the selected track's signal as a stereo source directly to only the rear Left Surround and Right Surround speakers in a 50/50 stereo balance.


Center + LFE, Stereo When selected for a stereo track, this option routes the track's left channel signal to the Center channel and the track's right channel signal to the LFE (subwoofer) channel discretely. If this option is selected for a track containing a Mono source file, the same signal is sent equally to both the Center and LFE channels. Note that this option is most useful with a stereo source file.

Center Only, Mono; FL Only, Mono; FR Only, Mono; Ls Only, Mono; Rs Only, Mono Sums the selected track's audio to a mono signal, and sends it all to the selected channel. This is the same as dragging the Panner Point directly onto one of the five main speakers in the Surround Panner.

Using the Surround Panner

The Surround Panner is an interactive control representing the audio field. You drag the Panner Point (white dot) to change the perceived sound source. As you move the Panner Point, the light blue Power Indicator lines coming from the speakers change length. The length of the lines indicates the power balance of your sound source coming from each of the five main channels. Additionally, a portion of the sphere appears dark blue to indicate the image of the sound field. That is, when seated in the center of the speakers (marked by the crosshairs), the blue area indicates where the listener perceives the sound to be coming from.

You can also drag the Panner Point outside the sound field directly on top of one of the five main speakers or on top of the LFE speaker. Once the Panner Point is in any one of these speaker locations, the audio from the currently selected track is summed to a mono signal and sent discretely to this one speaker channel. This is an easy way to send the complete track signal all to one channel.



Surround Panner options

A. Left Surround **B.** Front Left **C.** Center
D. Front Right **E.** Right Surround
F. Low Frequency Effects (Sub Bass) **G.** Panner Point


To use the Surround Panner:

In the Multichannel Encoder dialog box, drag the white Panner Point, which represents the location of the audio track in the sound field.

Automating the pan envelope

When you select Pan Envelopes, two envelope lines appear in the waveform display. The yellow envelope line controls the Left/Right balance and the green line controls the Front/Surround balance. These envelopes are interactive with the positioning of the Panner Point in the Surround Panner interface. It is possible to create dynamic panning over time by using these envelopes. (See [“Automating mixes with clip envelopes” on page 207.](#))

If you prefer to keep your track panned to a fixed point throughout the duration session, deselect Pan Envelopes. Deselecting this option removes the envelopes from the waveform display and lets you set the Panner Point to any static position you want. You can toggle the Pan Envelopes setting on and off and any envelope points you have created for this track are retained. Note that if Pan Envelopes is not selected, you can drag the Panner Point during playback and hear your static pan positioning in real time.



Pan Envelope Automation

To create a dynamic pan on a track:

- 1 In the Multichannel Encoder dialog box, select the box for a track in the Track List.
- 2 From the Panning Assignment menu, choose either “Surround Panner, Stereo Source” or “Surround Panner, Summed To Mono”.
- 3 Select Pan Envelopes, located above the right side of the waveform display. Two envelope lines appear in the waveform display. (Because the yellow line starts on top of the green line, you may see only the yellow line until you change the pan position.)
- 4 Click in the waveform display where you want to set a pan destination for the sound source. The vertical cursor moves to this time location.
- 5 Drag the Panner Point in the Surround Panner to the desired position in the sound field. Two handle points appear on the envelope lines within the waveform display and move as you position the Panner Point. (You can also click either of the envelope lines to create additional adjustable handles for shaping the envelope lines.)
- 6 To edit an envelope handle, drag it. The Panner Point moves in tandem to show you the relative position in the sound field during playback. To delete a handle, drag it up or down beyond the edge of the waveform display area.
- 7 To clear all envelope handle points and reset the track to flat envelopes, select Clear All, located below and to the right of the Pan Envelopes option.
- 8 To use spline curves for smoother transitions between points, select Splines.
- 9 Drag the playback cursor back to the start of the track, and select one of the Play buttons. Watch the Panner Point position, and listen for the dynamic pan setting you just created.

Adjusting volume levels

The Multichannel Encoder lets you adjust the subchannel level, center channel level, and track level.

To adjust the level:

Use any of the following options in the Multichannel Encoder dialog box:

Sub Channel Level Specifies the amplitude of the subchannel level to additionally send the track’s signal to the LFE channel. If the currently selected track is assigned to only the LFE channel, this option attenuates the amount of this track’s output sent to the LFE channel.

Note: *The Multichannel Encoder does not apply filtering to audio sent to the LFE channel, nor during preview, exporting, or encoding. Therefore, any low-pass filtering needed for your final LFE channel content should be applied to your audio within the Adobe Audition Multitrack, or on your exported .wav files.*



Use a Bass Management circuit in your monitoring setup to ensure that you hear the representative mix levels that might be reproduced in an end listener’s playback system.

Center Channel Level Determines the balance of the Front Left, Center, and Front Right channels when in the Surround Panner modes. When set to 100, the Center channel receives an equal percentage of signal as the Front Left and Front Right channels. The position of the Panner Point then determines the positional panning according to this Front Left, Center, Front Right balance ratio.

Track Level Controls the amplitude level of the currently selected track within the Multichannel mix in any selected Surround Panner mode.

Zooming into and out of the waveform display

There are several options for zooming in and out within the waveform display.


To zoom the waveform:

In the Multichannel Encoder dialog box, do any of the following:

- Place the mouse cursor over the time ruler that runs across the bottom of the waveform display, right-click and select a zoom option.
- Right-drag the desired zoom area on the time ruler. (To zoom back out again, right-click and choose Zoom Out or Zoom Full from the context menu.)
- Place the mouse pointer within the waveform display and turn the mouse wheel. This zooms into the time area directly beneath the mouse pointer. Turning the mouse wheel the other direction zooms back out.

Previewing the multichannel project

The Multichannel Encoder provides several preview options, including different types of playback controls and playback format options.



Preview controls

A. Transport controls **B.** Time indicator **C.** Output meters **D.** Preview device
E. Preview volume **F.** Preview device change button **G.** Master Level

Previewing a track or session in the Multitrack Encoder

You can choose from or adjust the following preview controls:

Go To Beginning ◀ Places the cursor at the start of the track.

Play Track Plays the currently selected track from the cursor location. Playback always plays to the end of the track, regardless of the current zoom level.

Play All Plays from the cursor location, playing the multichannel mix with all tracks that are checked in the Track List. Like Play Track, Playback always plays to the end of the session, regardless of the current zoom level.

Time indicator Located next to Play All, this indicator shows the time in the preview playback.

Preview Volume Controls the volume of the preview playback, without affecting the volume of the exported files. That is, it doesn't change the amplitude of the exported WAV or encoded WMA files that are created from the Multichannel Encoder, nor does it affect the levels measured by the 6-channel Output Meter. Use the Master Level slider for changing these.

Output meters (FL, FR, C, LFE, Ls, Rs) This set of six meters displays the output of each channel during preview. During Play Track, the meters display the output of only the selected track. During Play All, the meters display the output of the complete 5.1 mix. These levels are what the actual levels will be for your exported WAV or WMA files from the session. You can attenuate the overall 6-channel level by using the Master Level slider beneath the meters.

Master Level Sets the audible level of your preview playback. However, it is primarily intended to adjust the amplitude of the exported or encoded files. Use this slider and reference the meters to optimize the overall peak amplitude of the 5.1 channel mix so that none of the channels are clipping.

Preview Device, Format Displays the currently selected device to which Adobe Audition routes its 6-channel output. (See ["Setting the preview device and format" on page 233.](#))

Setting the preview device and format

The Preview Device, Format option lets you specify the device and format of the previewed audio. This option also displays the currently selected bit rate for preview playback. For information about device requirements, see ["Using the Multichannel Encoder" on page 226.](#)

To make changes to the device and bit selections:

- 1 Click the Change button to the right of the Preview Device, Format option.
- 2 From the Multichannel Output Device menu, specify the preview sound card.

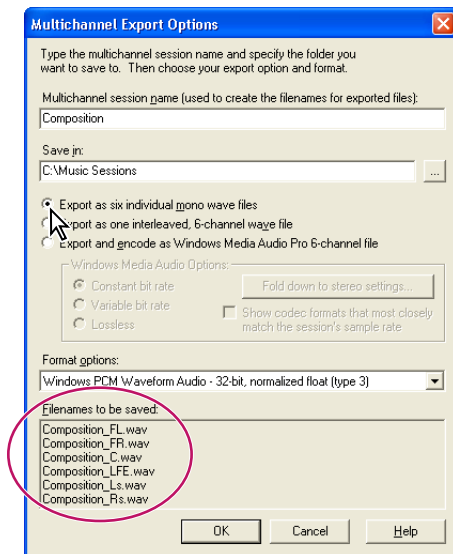
Note: Some sound cards that offer 5.1 playback, such as the Creative Labs Audigy, display only one device driver in the list. In this case, this is the device you should select because the sound card's driver will route the six channels of audio to the correct speakers. (If a multichannel device driver is not available from the sound card's manufacturer, you probably won't get a true surround-sound preview.) For sound cards that offer an interleaved multichannel driver, you should select this from the list. These driver types will accept the 6-audio input from Adobe Audition and automatically route it to the standard Microsoft 5.1 channel configuration.

- 3 From the Preview Format Selector menu, select the bit rate of the preview playback material that is sent to your sound card. If your session includes higher bit rate files, and if your sound card supports it, you can select a higher rate to the preview your session more accurately.
- 4 Set the Preview Buffer Size for the buffers used for Play Track and Play All. Larger buffer sizes enable a more stable preview playback, but increase the latency (that is, makes it take longer to play the result of changes made while previewing). If dropouts occur when you preview the audio, try increasing the buffer size.

Exporting surround-sound files

Adobe Audition includes the ability to encode directly to an interleaved 6-channel Windows Media 9 Pro (WMA) file or to export into two WAV formats. The Format Options field in the Multichannel Export Options dialog box indicates the currently selected format. The format is retained from your last used export option.

Note: To export and encode your project to a 6-channel WMA file, you must have Windows Media 9 runtime installed (it is installed by default as part of the Adobe Audition installation). If you have an earlier version installed, the Encode to WMA9 option will not be available. The latest Windows Media Updates are available on the Microsoft Web site.



The Multichannel Export Options dialog box

Exporting a multichannel session

Once you've completed mixing your multichannel project, you can export it to your desired file format.

To export your multichannel session:


- 1 Click Export at the bottom right of the Multichannel Encoder dialog box.
- 2 In the Multichannel Session Name text box, enter a name for your exported files. You can see the names for all the files that will be saved at the bottom of the dialog box in the Filenames To Be Saved area.
- 3 In the Save In text box, enter or navigate to the directory to which you want to save the files.
- 4 Select one of the following:

Export As Six Individual Mono Wave Files Creates standard Windows PCM .wav mono files that typically can be used by any Windows audio application. (For more information, see [“Windows PCM \(.wav\)” on page 247.](#))

Export As One Interleaved, 6-Channel Wave File Exports as Windows PCM .wav format, which allows a single file to contain multiple channels of audio. However, not all Windows audio applications can open or play WAV files that are not mono or stereo. (For more information, see [“Windows PCM \(.wav\)” on page 247.](#))

Note: *Interleaved files reflect the channel order used by Dolby Digital encoders. If you plan to use an encoding process with a different channel order, export the session as six individual files, instead.*

Export And Encode As Windows Media Audio Pro 6-Channel File Creates multi-channel WMA files that can be played by anyone who has Windows Media Player 9 or later, a multichannel output sound card, and a 5.1 speaker setup. (Media Player 9 requires Windows XP.)



Export options

5 If you select Export And Encode As Windows Media Audio Pro 6-Channel File, specify the following Windows Media Audio options:

- Constant Bit Rate (CBR) varies the quality level as needed to ensure that the bit rate stays the same. This method makes a consistently sized file, although the quality may not be as high as with Variable Bit Rate encoding.
- Variable Bit Rate (VBR) maintains the audio quality by varying the bit rate depending on the complexity of the audio passage being encoded. This method can maintain higher quality audio in the file, although the file size is not as predictable as with Constant Bit Rate encoding.
- Lossless compresses to a smaller file size than WAV, but results in no fidelity loss.
- Fold Down To Stereo Settings folds down the 6-channel playback to a stereo playback on a non-Windows XP system or a system without a 5.1 playback setup. Specify the attenuation parameters to control how the levels of the Center, Surround, and LFE channels get mixed down with the front stereo channels and played back on a stereo output system. The defaults are usually good settings for most files, but you can enter any value in any of these three fields between 0 and -144 dB, as desired.
- Show Codec Formats That Most Closely Match The Session's Sample Rate limits the list of selectable WMA kbps options to those that are the same sample and bit rate as the multitrack session's files.

Exporting to a mastering or duplication service

If your project is to be sent out to a mastering, duplication, or other outside service with the intention of being encoded into other specific surround or media formats, check with the recipient as to the format specifics.

Channel ordering differs between surround formats, as do the crossover frequency points. For example, Digital Theater System (DTS) typically employs a crossover of 80 Hz, meaning that all frequency content of your channels lower than 80 Hz can be routed to a subwoofer, and all frequency content greater than 80 can be sent to the main channels.

This differs from the Dolby Digital system that utilizes a crossover point of 120 Hz. Some systems also employ a boost of 10 dB for the LFE channel, automatically assuming your LFE content will be approximately this much lower in power than the main channels.

Therefore, these components should be accounted for in your mix before you deliver your master files to the recipient. It is best to inquire with the recipient about all such requirements to ensure that the audience will hear your project the same way you are hearing it on your monitoring system.




Saving, Exporting, and Closing Files

Saving audio files

When working with audio files in Edit View, you can save your audio in a variety of common file formats. The format you choose depends on how you plan to use the file. For more information on supported file formats, see [“Choosing an audio file format” on page 240](#).

When choosing a file format, keep in mind that different formats allow different information to be stored with the file. As a result, saving a file in a format different from its original format might cause some information to be discarded.

To save an audio file in Edit View:

- 1 Do one of the following:
 - Choose File > Save to save changes you made to the current file. Alternatively, click the Save button  in the File toolbar.
 - Choose File > Save As to save changes to a different file. Alternatively, click the Save As button  in the File toolbar.
 - Choose File > Save Copy As to save an identical copy of the file while leaving the original file active.
 - Choose File > Save Selection to save the currently selected audio to a new file. Alternatively, click the Save Selection button  in the File toolbar. This command is useful for saving small segments of a larger file. For example, you can use it to break up a long recording into smaller, more manageable tracks.
 - Choose File > Save All to save all open files.
- 2 Choose a location for the file, type a filename, and choose a file format.
- 3 Depending on the format you choose, additional options might be available. To view format-specific options, click Options. For more information on format-specific options, see [“Choosing an audio file format” on page 240](#).
- 4 Select Save Extra Non-Audio Information to save header fields containing file information and cue marks in the file. In addition, if you save a .wav file, this option stores the pathname to the original session file, effectively linking related session and mixdown files for Adobe Premiere and Adobe After Effects users. For more information, see [“Working with Adobe Premiere Pro and After Effects” on page 222](#).



If you plan to burn the file to CD by using another program, you should deselect this option. Some CD recording applications interpret non-audio information incorrectly and place an unpleasant burst of noise at the beginning of each track.

- 5 Click Save.



Saving and exporting sessions



When you edit a session in Multitrack View, it's a good practice to save the session file frequently. After you create a mix, you can export the session to a variety of audio and video formats.

Saving sessions

The most important thing to remember about session (.ses) files in Adobe Audition is that they contain no audio data themselves. Instead, a session file is a small file that points to other audio files on your hard drive. The session file keeps track of where the audio files are stored on the hard drive, each file's location and duration within the session, what envelopes and effects are applied to the tracks, and so on.

A session file is useless without the audio files that it points to, so it's important to keep your files organized. The best way to stay organized is to keep all session-related files in the same folder. Adobe Audition makes organizing files easy by providing an option to save a copy of every file used in a session into the same folder as the session file. This option ensures that all files for a session are in one place.

To save a session:

- 1 Do one of the following:
 - Choose File > Save Session to save changes to the current session file. Alternatively, click the Save button  in the Multitrack File toolbar.
 - Choose File > Save Session As to save changes to a different session file. Alternatively, click the Save As button  in the Multitrack File toolbar.
 - Choose File > Save All to save all open sessions.
- 2 Choose a location for the file, and type a filename.
- 3 Select Save Copies Of All Associated Files to save a copy of every file used in a session into the same folder as the session file. It is highly recommended that you select this option.

If you want to save the associated files in a different format, click Options, select Save All Copies In This Format, and select a format from the list. To view options for the selected format, click Format Properties. For more information on format-specific options, see [“Choosing an audio file format” on page 240](#).
- 4 Click Save.

To convert the sample rate of a session:

- 1 Choose File > Save Session As, choose a location for the file, and type a filename.
- 2 Select Save Copies Of All Associated Files, and click Options.
- 3 Select Convert Sample Rate, and select a sample rate.
- 4 To set dithering and other conversion options, click Conversion Properties. For more information on conversion options, see [“Converting the sample type” on page 93](#).
- 5 Click Save.

Exporting mixes to audio

After you finish mixing a session, you can export it in a variety of common audio file formats. When you use the File > Export > Audio command, everything in the session is exported to an audio file. If you want to export only specific waves, use the Edit > Mix Down To File command instead. (See [“Mixing down ReWire tracks and specific audio clips” on page 213.](#))

To export a mix to an audio file:

- 1 Do one of the following:
 - To export part of a session, select the desired area in the track display.
 - To export an entire session, deselect everything in the track display. (If necessary, click the track display to reveal the current-time indicator.)
- 2 Choose File > Export > Audio.
- 3 Choose a location for the file, type a filename, and choose a file format.
- 4 Depending on the format you choose, additional options may be available. To view format-specific options, click Options. For more information on format-specific options, see [“Choosing an audio file format” on page 240.](#)
- 5 Select Save Extra Non-Audio Information to save header fields containing file information and cue marks in the file. In addition, if you save a .wav file, this option stores the pathname to the original session file, effectively linking related session and mixdown files for Adobe Premiere and Adobe After Effects users. For more information, see [“Working with Adobe Premiere Pro and After Effects” on page 222.](#)



If you plan to burn the file to CD by using another program, you should deselect this option. Some CD recording applications interpret non-audio information incorrectly and place an unpleasant burst of noise at the beginning of each track.

- 6 Click Save.

Exporting mixes to video

If a session includes an .avi video file, you can mix down the session and make it an audio track for the video.

Note: While Adobe Audition can open other types of video files to get at their audio tracks, the ability to save back the audio track works only with .avi files.

To export a mix to a video file:

- 1 Choose File > Export > Video.
- 2 Choose a location for the file, and type a filename.
- 3 To assign a codec for compressing the audio in the file, click Options. Choose a codec from the pop-up menu, and click OK.
- 4 Select Save Extra Non-Audio Information to save header fields containing file information and cue marks in the file.
- 5 Click Save.

Closing files

Adobe Audition provides several commands for closing files.

To close the current audio file in Edit View:

Choose File > Close.

To close a session file in Multitrack View:

Do one of the following:

- Choose File > Close Session to close the current session file but leave related media files open.
- Choose File Close Session And Its Media to close the current session file and all related media files.

To close all files not related to the current session:

Choose File > Close Only Non-Session Media Files.

To close all open files:

Choose File > Close All.

Choosing an audio file format

Adobe Audition lets you open and save files in the formats described in this section. In most cases, you should save uncompressed audio in Windows PCM format, and you should save compressed audio in either mp3PRO® format or Windows Media Audio format. You'll need to use other formats only in special situations.

Some formats provide options for saving audio data. Click Options in the Save As dialog box to access these options.

Note: If you want to save files in a format that's not listed here, you may be able to do so by using an ACM Waveform codec. For more information, see [“ACM Waveform \(.wav\)” on page 241](#).

64-bit doubles (RAW) (.dbl)

This format uses 8-byte doubles in binary form—8 bytes per sample mono, or 16 bytes per sample stereo interleaved. The 64-bit doubles format has no header—it's purely audio data, just like the raw PCM format.

8-bit signed (.sam)

This format is popular for building MOD files, since audio in MOD files is 8-bit signed. Many MOD editors allow samples to be inserted from or exported to files in this format. Files with the .sam extension contain 8-bit signed raw data, and by default, they have no headers. The sample rate starts off as 22050 Hz, but you can change the sample rate after you open the file by choosing Edit > Adjust Sample Rate.

A/mu-Law Wave (.wav)

The A-Law and mu-Law WAV formats (CCITT standard G.711) are common in telephony applications. These encoding formats compress the original 16-bit audio to 8-bit audio (for a 2:1 compression ratio) with a dynamic range of about 13-bits (78 dB). While A-Law and mu-Law encoded waveforms have a higher signal-to-noise ratio than 8-bit PCM, they also have a bit more distortion than the original 16-bit audio. Still, the quality is higher than you would get with some 4-bit ADPCM formats.

Note: Files saved in this format expand automatically to 16-bits when opened, so you shouldn't save 8-bit files in this format.

Options Choose from the following:

- A-Law 8-bit is a slight variation of the standard mu-Law format and is found in European systems.
- mu-Law 8-bit is the international standard telecommunications encoding format and is the default option.

ACM Waveform (.wav)

Microsoft ACM (Audio Compression Manager) is part of all 32-bit versions of Windows. Adobe Audition supports the ACM driver, which enables you to open and save files in a variety of formats other than those directly supported by Adobe Audition.

Some of these formats come as a standard part of Windows, while others are provided by third-parties. You may acquire ACM formats when you install other software.

To save a file in an alternate format by using the ACM driver, choose File > Save As, choose ACM Waveform as the file format, and click Options. You can select from among various quality levels, and each level will give you different options for formats and attributes.

Note: The ACM driver you want to use might require that the file be in a specific format before saving it. For example, if you want to save a file in the DSP Group TrueSpeech format, you should first use the Edit > Convert Sample Type command to convert the file to 8 KHz, mono, 16-bit, because that is the only format that the TrueSpeech ACM driver supports. For more information on any particular ACM driver, contact the creator of the format (such as DSP Group for TrueSpeech, or CCITT for the various CCITT formats) or the manufacturer of the hardware that uses the format in question.

Amiga IFF-8SVX (.iff, .svx)

The Amiga IFF-8SVX format is an 8-bit mono format from the Commodore Amiga computer.

Options Choose from the following:

- Data Formatted As saves the audio file in uncompressed 8-bit Signed format (the default setting) or in the compressed 4-bit Fibonacci Delta Encoded format.
- Dithering From 16-bit specifies a type of dithering for 16-bit files: Triangular Dither, Shaped Gaussian Dither, Noise Shaping A, or Noise Shaping B. No Dithering is the default. For more information on types of dithering, see [“Changing the bit depth” on page 96](#).

Apple AIFF (.aif, .snd)

AIFF is Apple's standard wave file format. AIFF supports mono or stereo files, 16-bit or 8-bit resolution, and a wide range of sample rates. Adobe Audition supports only the PCM-encoded portion of the data, even though this format (like Windows WAV) can contain any one of various data formats.

AIFF is a good choice for Windows/Mac OS cross-platform compatibility. Before you open AIFF files in Adobe Audition, add the .aif or .snd extension to the file and open it by using the Apple AIFF file filter. When you transfer an AIFF file to a Macintosh, you can add the four character code "AIFF" in the file's resource fork to have it recognized. (The Macintosh identifies a file through its "resource," which is removed when a file is opened on a Windows computer. However, many Mac OS applications that support AIFF can recognize the PCM data without this identifier.)

ASCII Text Data (.txt)

Audio data can be read to or written from files in a standard text format, with each sample separated by a carriage return, and channels separated by a tab character. An optional header can be placed before the data. If no header text exists, then the data is assumed to be 16-bit signed decimal integers. The header is formatted as a KEYWORD: value with the keywords being SAMPLES, BITSPERSAMPLE, CHANNELS, SAMPLERATE, and NORMALIZED. The values for NORMALIZED are either TRUE or FALSE. For example,

```
SAMPLES: 1582
BITSPERSAMPLE: 16
CHANNELS: 2
SAMPLERATE: 22050
NORMALIZED: FALSE
164 <tab> -1372
492 <tab> -876
```

Options Choose any of the following:

- Include Format Header places a header before the data.
- Normalized Data normalizes the data between -1.0 and 1.0.

Audition Loop (.cel)

This format produces compressed Adobe Audition loop files, which are essentially .mp3 files with a .cel extension. Each .cel file has a header that contains loop information, such as the number of beats, tempo, key, and stretch method.



You can also save loops in uncompressed formats, such as Windows PCM.

The .cel format avoids a potential problem with .mp3 files. During encoding, a very small amount of silence is added to the beginning, end, or both of an .mp3 file. The silence is very short—often only a few samples long. When you work with a loop, though, it's enough to throw off the entire loop.

As it saves a .cel file, Adobe Audition calculates how much silence will be added to the .mp3 file and writes this information into the .cel header. Then, when Adobe Audition opens a .cel file, it reads this information and automatically removes the silence from the file so that it loops smoothly.

The options for Audition Loop format are identical to those for mp3PRO®. For more information, see [“mp3PRO® \(.mp3\)” on page 244](#).

Creative Sound Blaster (.voc)

This format is for Sound Blaster and Sound Blaster Pro voice files. Adobe Audition supports both the older and newer formats. The older format supports only 8-bit audio, mono to 44.1 kHz and stereo to 22 kHz. The newer format supports both 8- and 16-bit audio.

Files in this format can contain information for looping and silence. If a file contains loops and silence blocks, they expand when you open the file.

Options Choose one of the following:

- Old Style saves audio as an 8-bit .voc file that can be played on any Sound Blaster card.
- New Style saves audio to the newer format that supports both 8- and 16-bit audio.

Dialogic ADPCM (.vox)

The Dialogic ADPCM format is used in telephony applications, and it's optimized for low sample rate voice. It supports only mono 16-bit audio, and like other ADPCM formats, it compresses the audio data to 4 bits/sample (4:1). This format has no header, so Adobe Audition assumes any .vox file is in Dialogic ADPCM format.

Note: Take note of the sample rate of the audio before saving it, as you need to enter it upon reopening the file.

DiamondWare Digitized (.dwd)

This format is used by DiamondWare Sound Toolkit, a programmer's library that lets you quickly and easily add high-quality interactive audio to games and multimedia applications. It supports both mono and stereo files at a variety of resolutions and sample rates.

DVI/IMA ADPCM (.wav)

The International Multimedia Association (IMA) flavor of ADPCM compresses 16-bit data to 4 bits/sample (4:1) by using a different (faster) method than Microsoft ADPCM. It has different distortion characteristics, which can produce either better or worse results depending on the sample being compressed. As with Microsoft ADPCM, use this format with 16-bit rather than 8-bit files. This compression scheme can be a good alternative to MPEG; it provides reasonably fast decoding of 4:1 compression, and it degrades sample quality only slightly.

Options Choose from the following:

- 2 bits/sample, 8:1 produces files with the highest compression ratio (8:1) but with the lowest number of bits. Select this option if smaller file size is more important than audio quality. Keep in mind that this compression rate is less compatible than the standard 4-bit and is supported on fewer systems.
- 3 bits/sample, 5.3:1 produces higher quality than the 2 bits option, but the quality isn't quite as good as with the 4 bits and 5 bits options. Some systems might have problems playing back files with this compression rate, especially stereo files.

- 4 bits/sample, 4:1 produces 4-bit files at a compression ratio of 4:1. This option is the default.
- 5 bits/sample, 3.2:1 produces files with the highest quality, since more bits and a lower compression ratio are used. However, this compression rate is less compatible than the standard 4-bit.

Microsoft ADPCM (.wav)

The Microsoft ADPCM format provides 4:1 compression. Files saved in this format expand automatically to 16-bits when opened, regardless of their original resolution. For this reason, use this format with 16-bit rather than 8-bit files.

Options Choose from the following:

- Single Pass (Lower Quality) compresses files in a single pass. Use this option if you're in a hurry. However, the quality is lower than if you use the Multiple Pass option. The time taken to read an ADPCM-compressed file is the same no matter which option you use.
- Multiple Pass (Higher Quality) compresses files in multiple passes, providing better quality. This setting is the default.
- Block Size offers three size options, each with a different compression ratio and quality level: Large (Default Quality), with a compression ratio of 3.98:1; Medium (Good Quality), with a compression ratio of 3.81:1; and Small (High Quality), with a compression ratio of 3.25:1.

mp3PRO® (.mp3)

The mp3PRO filter enables Adobe Audition to directly encode and decode .mp3 files. When you save a file to mp3 format, the audio is encoded and compressed according to the options you select. When you open an .mp3 file, the audio converts into the uncompressed internal format of Adobe Audition. As a result, you can save an .mp3 file in any format.



Avoid compressing the same audio to mp3 more than once. Opening and resaving an .mp3 file causes it to be recompressed, so any artifacts from the compressing process become more pronounced.

MP3/mp3PRO® Encoder Options dialog box contains two sets of options: basic options for choosing an encoding method and more advanced options. To view the advanced options, click Advanced. To view only the basic options, click Simple.

Basic options Choose from the following:

- CBR (constant bit rate) encodes the same bit rate throughout the entire file. This method is the most common and the most predictable for bandwidth and file size.
- VBR (variable bit rate) encodes higher bit rates for more complex material and lower bit rates for simpler material. While it depends on the source material, VBR-encoded .mp3 files generally tend to be smaller than CBR-encoded .mp3 files. Use the menu below the VBR option to choose a quality level from 10 (lowest quality but smaller file) to 100 (highest quality but larger file). Some mp3 players don't support VBR-encoded files. For maximum compatibility, select CBR.
- MP3 encodes the file to mp3, but without the PRO data.

- mp3PRO® encodes the file mp3PRO. The PRO data helps re-create high frequencies in the compressed file, especially at low bit rates. An mp3PRO file can still be played back by an mp3 player that doesn't support the PRO data, but the quality may be lower than for a standard mp3 file of that bit rate. For example, a 64 Kbps mp3PRO file sounds more like a 112 Kbps or 128 Kbps mp3 file if the player supports mp3PRO, but it sounds like a 64 Kbps mp3 file (or worse) if the player doesn't support mp3PRO.

Advanced options Click Advanced to choose from the following:

- Maximum Bandwidth (available only if MP3 is selected) specifies the highest frequency that will be encoded. Lower bandwidths help eliminate tinkly and phase-like effects but at the expense of reducing the higher frequencies.
- CBR Bitrate (available only if CBR and MP3 are selected) specifies the bit rate for CBR encoding. The higher the number, the larger the file, but the better the quality. Valid values range from 20 Kbps to 320 Kbps.
- Sample Rate (available only if CBR and MP3 are selected) specifies the sample rate of the destination file. (The decoder will also use this rate.) Keep in mind that not all sample rates are valid for a particular bit rate.
- VBR Quality (available only if VBR is selected) specifies the quality for VBR encoding. The higher the number, the larger the file, but the better the quality. Valid values range from 1 to 100.
- Low Complexity Stereo (available only if CBR and mp3PRO® are selected) encodes the audio as mono, with information on how to reconstruct the stereo signal at playback. A non-PRO decoder plays back only mono, but a PRO decoder plays back stereo. The stereo image is different from the original audio, but it most often sounds better than its mono counterpart.
- Codec provides three codec options. Depending on the type of audio, one codec might do a better job than the others. Experiment to see which one is best for your project. Current-Best Quality is an extremely fast algorithm, and it generally does a very good job at lower bit rates, as well as giving more high-frequency detail without unwanted artifacts. Unless you have a specific reason not to, use this setting. Legacy-Medium Quality (Fast) uses a different model for encoding and can be more complete at bit rates above 160 Kbps. Legacy-High Quality (Slow) takes longer to encode, but the quality is higher than that of the Medium Quality option.
- Allow Mid-Side Joint Stereo (for mid bit rates) combines the left and right channels by using a Mid-Side method when encoding middle quality bit rates and below. This option preserves surround-sound information by saving the common audio in one channel while the difference between the channels is saved in the other.
- Allow Intensity Joint Stereo (for low bit rates) allows combining of the left and right channels for files encoded at low bit rates. Some frequencies are saved as mono and placed in the stereo field based on the intensity of the sound.

Note: Don't use this option if the stereo audio contains surround-encoded material.

- Allow Narrowing Of Stereo Image uses more data to represent a wider stereo image. This option allows the encoder to narrow the image in some parts in order to make the overall audio quality better.
- Set 'Private' Bit sets the Private bit for each MPEG frame.
- Set 'Copyright' Bit sets the Copyrighted bit on the .mp3 file.
- Set 'Original' Bit sets the Original Copy bit, which designates that the .mp3 file is on its original media.
- Padding specifies a padding option. ISO Padding is the default, but you can choose a different setting if the decoder needs no padding or always needs padding.
- Set All Decoding To 32-Bit determines how .mp3 files are opened in Adobe Audition. Selecting this option forces Adobe Audition to upsample non-32-bit .mp3 files to 32-bit. Deselecting this option allows .mp3 files to be opened with the original bit depth intact.
- Encode Stereo As Dual Channel encodes two audio channels with independent contents within one bitstream.
- Write CRC Checksums adds CRC checksums to the audio stream so that content can be verified for any errors when decoded.

NeXT/Sun (.au, .snd)

The NeXT/Sun format is standard on NeXT and Sun computers, and it has many data types. Adobe Audition supports the CCITT A-Law, mu-Law, G.721 ADPCM, and linear PCM data variants. Like Windows PCM and AIFF, this format can support mono or stereo, 16- or 8-bit, and a wide range of sample rates when saved as linear PCM.

The NeXT/Sun format is most commonly used for compressing 16-bit data to 8-bit mu-law data. AU is used quite extensively on the Web and in Java applications and applets.

Options Choose from the following:

- mu-Law 8-bit uses the mu-law 8-bit format to compress the file.
- A-Law 8-bit uses the A-law 8-bit format to compress the file.
- G.721 ADPCM 4-bit applies the standard CCITT G.721 compression to the file (ADPCM at 32Kbps).
- Linear PCM saves the file as uncompressed, linear PCM (Pulse Code Modulation).

SampleVision (.smp)

The SampleVision format is native to Turtle Beach's SampleVision program. This format supports only mono, 16-bit audio. If a file is in a different format, Adobe Audition prompts you to convert it before saving it.

This format also supports loop points, which you can edit in the Cue List window. The Label of the cue must be in the format `LOOP n, m` where "n" is the loop number from 1 to 8, and "m" is the mode (0 = no looping, 1 = forward loop, 2 = forward/back loop). In the Play List window, you can enter the number of times to loop the cue range.

Windows Media Audio (.wma)

The WMA format utilizes a perceptual compression scheme and lets you select from three different encoding options:

- Constant Bit Rate Encoding varies the quality level as needed to ensure that the bit rate stays the same. This method makes a consistently sized file, although the quality may not be as high as with Variable Bit Rate encoding.
- Variable Bit Rate Encoding maintains the audio quality by varying the bit rate depending on the complexity of the audio passage being encoded. This method can maintain higher quality audio in the file, although the file size is not as predictable as with Constant Bit Rate encoding.
- Mathematically Lossless Encoding compresses to a smaller file size than WAV, but results in no fidelity loss.

After you select an encoding option, you can set the desired quality. Just as with stereo WMA files, the higher quality setting you select, the larger the file size, and vice versa.

Windows PCM (.wav)

The Microsoft Windows PCM format supports both mono and stereo files at a variety of resolutions and sample rates. It follows the RIFF (Resource Information File Format) specification and allows for extra user-information to be embedded and saved with the file. The WAV format reproduces digital audio by using PCM (Pulse Code Modulation)—PCM doesn't require compression and is considered a lossless format.

Options The following options are available for 32-bit files; no options are available for 8- or 16-bit files:

- 32-bit Normalized Float (type 3) – Default is the internal format for Adobe Audition and the standard floating point format for type 3 .wav files. Values are normalized to the range of +/-1.0, and although values beyond this range are saved, clipping may occur in some programs that read them back in. (Adobe Audition won't clip audio but will instead read the same value back if it's beyond this range.)
- 24-bit Packed Int (type 1, 24-bit) saves straight 24-bit integers so any data beyond the bounds is clipped. The .wav BitsPerSample is set to 24 and BlockAlign is set to 3 bytes per channel.
- 24-bit Packed Int (type 1, 20-bit) saves straight 24-bit integers so any data beyond the bounds is clipped. The .wav BitsPerSample is set to 20 and BlockAlign is set to 3 bytes per channel. The extra 4 bits are actually the remaining valid bits when saving, and they are used when reading (thus still giving 24-bit accuracy if those bits were actually present when writing). Applications either fill those last 4 bits with zeros or with actual data; analog/digital converters that generate 20 bits of valid data automatically set the remaining 4 bits to zero. Any type 1 format with BlockAlign set to 3 bytes per channel is assumed to be packed integers, and a BitsPerSample value between 17 and 24 will read in all 24 bits and assume the remaining bits are either accurate or set to zero.

- 32-bit 24.0 Float (type 1, 24-bit) – Non-Standard saves full 32-bit floats (ranging from ± 8 million), but the .wav BitsPerSample is set to 24 while BlockAlign is still set to 4 bytes per channel.
- 16.8 float – Obsolete/Compatibility is the internal format used by Adobe Audition 1.0. Floating point values range from ± 32768.0 , but larger and smaller values are valid and aren't clipped since the floating point exponent is saved as well. The .wav BitsPerSample is set to 32 and BlockAlign is set to 4 bytes per channel.
- Enable Dithering dithers 32-bit files when they are saved to a PCM format (20-bit, 24-bit, or 32-bit). This option is available only for a 32-bit file that you select to save to a nonfloating-point type format. It applies a Triangular dither with a 1.0 depth 1.0 and no noise shaping. If you wish to apply a noise-shaped dither, use the Edit > Convert Sample Type command to dither the audio first, and then save the file without dithering enabled in the file format options.

PCM Raw Data (.pcm) (.raw)

This format is simply the PCM dump of all data for the wave. No header information is contained in the file. For this reason, you must select the sample rate, resolution, and number of channels upon opening the file.

By opening audio data as PCM, you can interpret almost any audio file format—but make sure that you have some idea about the sample rate, number of channels, and so on. You can also interpret the data as A-law or mu-law compressed. When you guess at these parameters upon opening a file, it may sound incorrect (depending on which parameters are wrong). Once the file is opened and sounds fine, you may hear clicks at the start or end of the waveform, or sometimes throughout. These clicks are various header information being interpreted as waveform material. Just cut these out, and you've read in a wave in an unknown format.

Options Choose from the following:

- Data Formatted As specifies the format of the saved data.
- When Opening, Offset Input Data By specifies the number of bytes by which to offset the input data.
- Create .DAT Header File On Save writes a header to a separate .dat file to make reopening the file easier.

Scripting and Batch Processing

About scripting and batch processing

Adobe Audition scripts let you save a series of actions such as copying data or applying an effect, so you can perform those actions again with the click of a button. Scripts are simple text files that are similar to macros; Adobe Audition stores the exact actions of your mouse and any tweaking of parameters, so you can repeat them in the same sequence when you run the script.

For example, suppose you have a combination of effects with particular settings (an EQ setting, a Hall reverb, and so on) that you want to apply often and in combination to achieve a certain sound. You can record these steps, along with effects' specific settings, and then apply them at any time simply by calling the script.

Batch processing cue ranges

You can use the Batch feature in the Cue List to add silence between cues and save the audio between cues to new files. For more information about cues, the Cue List, and cue ranges, see [“Working with cues” on page 82](#).

To batch process cues:

- 1 Choose Window > Cue List.
- 2 Select one or more cues in the Cue List dialog box. At least one of the cues you select must be a range.
- 3 Click Batch at the bottom of the dialog box.
- 4 Set the following options as desired, and click OK:

Set Amount Of Silence Adds silence between cue points in the current waveform. Enter the number of seconds of silence you want in the Add Silence Before and Add Silence After text boxes.

Save To Files Splits the audio between cue points in the active waveform to new files.

Use Cue Label As Filename Uses the name of the cue as the prefix for the filename.

Filename Prefix Specifies the prefix for the filename (such as “phrase”). Adobe Audition automatically adds numbers after the prefix (phrase02, phrase03, and so on) in addition to the correct extension based on the output format you specify.

Seq. Start Specifies the number to begin with when appending numbers to the filename prefix.

Destination Folder Specifies the folder in which you want Adobe Audition to place new “split” files. Click Browse to open the Choose Destination Folder window and locate a different folder.

Output Format Sets the output format. Depending on the format, Options is available. Click Options to select options for that format.



Normalizing groups of files

When you normalize a waveform, the loudest part of the waveform is set to a specified amplitude, thereby raising or lowering all other parts of the same waveform by the same amount. Group Waveform Normalize lets you normalize the volume of multiple open waveforms by using a three-screen batch process. If the volume is raised as part of the normalization process, Adobe Audition can apply limiting to prevent clipping.

If you're getting ready to master an audio CD, using Group Waveform Normalize is a great way to make sure that all tracks on the CD have a consistent volume.

To normalize a group of files:

- 1 Choose Edit > Group Waveform Normalize.
- 2 Select the open waveforms you want to normalize. Click to select a single file, Shift-click to select contiguous files, Ctrl-click to select noncontiguous files, and drag to select a group of files.
- 3 Click the Analyze Loudness tab, and then click Scan For Statistical Information to display amplitude statistics for each waveform. Double-click a file in this list to see more detailed statistics, including an RMS histogram and a clipping profile. (See [“Understanding statistics on the Analyze Loudness tab” on page 250.](#))
- 4 Click the Normalize tab, and specify how you want to normalize the waveforms. (See [“Setting options on the Normalize tab” on page 251.](#))
- 5 Click Run Normalize.

Understanding statistics on the Analyze Loudness tab

When you click Scan For Statistical Information, the Analyze Loudness tab displays the following information:

Eq-Loud Is the final loudness value with an equal-loudness equalization curve that takes into account frequencies to which the human ear is most sensitive. If you select the Use Equal Loudness Contour option in the Normalize tab, this value determines how much to amplify the audio to normalize it.

Loud Is the final loudness value without equal-loudness equalization. If you don't select the Use Equal Loudness Contour option in the Normalize tab, this value determines how much to amplify the audio to normalize it.

Max Is the maximum RMS (Root-Mean-Square) amplitude present. This value is based on a full-scale sine wave being 0 dB, and it conforms to the width specified in the Advanced section of the Normalize tab.

Avg Is the average RMS of the entire waveform. This value isn't used for normalization.

% Clip Is the percentage of the waveform that would be clipped as a result of normalization. Clipping won't occur if limiting (in which loud passages are decreased in volume) is used; instead, the louder portions of audio are limited to prevent clipping. In general, avoid values higher than 5% to prevent audible artifacts from occurring in the louder portions of audio.

Reset Clears all of the normalization statistics for the files in the list.



Double-click a file in this list to see more detailed statistics, including a complete RMS histogram, which shows the relative amounts of audio at each loudness level, and a clipping profile, which shows how much clipping will occur for each decibel of amplification.

Setting options on the Normalize tab

Use the following options in the Normalize tab to specify how you want to normalize the waveforms:

Normalization Specifies whether to normalize to an average level or a specific level you enter in decibels.

Note: *The Normalization option doesn't use percentages, unlike the Normalize effect, because it is RMS-based rather than peak-based.*

Use Equal Loudness Contour Applies an equal loudness contour, where the middle frequencies are most important. Because the human ear is much more sensitive to frequencies between 2 kHz and 4 kHz, two different pieces of audio with the same RMS amplitude but with different frequencies will have different apparent volumes. Select this option to ensure that audio has the same perceived loudness, regardless of what frequencies are present.

Out of Band Peaks Determines how Adobe Audition handles out-of-band peaks. When you amplify audio, the audio samples may extend beyond the clipping point. If out-of-band peaks occur, you can choose to just let it clip the waveform (and cause distortion), or you can apply limiting to those areas so the audio doesn't clip the waveform (a common practice for TV commercials so they sound louder).

- No Limiting (Clip) prevents limiting, so clipped (distorted) audio might occur.
- Use Limiting applies the Hard Limiter, if needed, to keep out-of-band peaks from being clipped. This options provides two additional options: Look Ahead Time and Release Time.
- Lookahead Time specifies the number of milliseconds generally needed to attenuate audio before reaching the loudest peak.

Note: *If this value is too small, audible distortion might occur. Make sure that the value is at least 5 milliseconds.*

- Release Time specifies the number of milliseconds needed for the attenuation to rebound 12 dB (or roughly the time needed for audio to resume normal volume if an extremely loud peak is encountered).

Note: *A setting of 200 milliseconds works well to preserve low bass frequencies. If the setting is too high, audio may stay quiet and not resume normal levels for a while.*


Statistics RMS Width Specifies the length of the audio selection to use for calculating the RMS (Root-Mean-Square) minimum and maximum values.

Batch processing files

The Batch Processing dialog box in Adobe Audition enables you to run a single script repeatedly over a group (batch) of source files.

Note: *For a script to run on a batch of files, you must record it in Script Works On Current Wave mode. That is, before you record the script, a waveform must be open with no selection made.*

In addition, the Batch Processing dialog box enables you to change multiple waveforms from one audio format to another (such as from WAV to MP3). For more information, see [“Converting the sample type” on page 93](#).



The Files tab of the Batch Processing dialog box

To batch process files:

- 1 In Edit View, choose File > Batch Processing. The Batch Processing dialog box appears with the Files tab displayed.
- 2 Click Add Files to open the Please Choose The Source Files dialog box, and select one or more files:
 - Hold down Ctrl or Shift to select noncontiguous or contiguous files, respectively.
 - Click Remove to delete highlighted files from the list.
 - Click Remove All to delete all files from the list.
 - Click Hide Path to display the name of the file without its full path.
 - Click Open Raw PCM As to select the desired Sample Rate, Channels, Resolution, and other properties. Use this option only when converting Raw PCM files.
- 3 Click the Run Script tab at the bottom of the Batch Processing dialog box.
- 4 Select Run A Script. Then, click Browse to locate and select a script collection (*.scp) file, and click Open.
- 5 Choose a script from the Script menu. The only scripts that you can use for batch processing (and the only ones that appear in the list) are those that were recorded in Script Works On Current Wave mode. (See [“Working with scripts” on page 254](#).)
- 6 Click the Resample tab.
- 7 Select Conversion Settings to change each waveform’s sample properties to a common set of values. Then, click Change Destination Format to specify the values. (See [“Converting the sample type” on page 93](#).) If you don’t select Conversion Settings, the sample properties for the destination files are the same as those for the source files.
- 8 Click the New Format tab.
- 9 From the Output menu, choose a format for the destination files.

10 Click Format Properties to display options for the destination format.

Note: *Sample Format Types* lists the sample properties of the waveforms that are to be converted. If more than one entry is listed, you might have to select different properties for each, depending on the destination format. For instance, a 22 kHz mono waveform might need different encoding options than a 44 kHz stereo file.

11 Click the Destination tab.

12 Select a destination folder, specify how files are renamed by setting the following options, and then click Run Batch:

Same As File's Source Folder Saves modified files in the same folder as the file's source file.

Other Folder Specifies the folder in which to save modified files. Click Browse to locate a folder.

Overwrite Existing Files Saves existing files with a new name.

Delete Source File If Converted OK Deletes source files after they are converted successfully.

Remove From Source List If Converted OK Removes filenames from the source list after the files are converted successfully.

Output Filename Template Specifies how files are renamed. By default, the first part of the filename remains the same, and the extension changes to match the chosen output format. Alternatively, you can type a different extension, and you can set up conditions for how files are renamed by using question marks and asterisks:

- A question mark ("?") signifies that a character doesn't change.
- An asterisk ("*") denotes the original filename or extension.

Here are some examples of how filenames can be renamed:

Original Name	Output Filename Template Name	Resulting Filename
zippy.aif	*.wav	zippy.wav
toads.pcm	q*.voc	qtoads.voc
funny.mp3	b??????.*	bunny.mp3
biglong.au	????.au	bigl.au
bart.wav	*x.wav	bartx.wav

Working with scripts

Adobe Audition lets you create three types of scripts, depending on the software's state when you record the script:

- Scripts that start from scratch. These scripts start with no waveform opened, and their first command is File > New.
- Scripts that work on the currently open waveform. These scripts operate on an entire waveform. They require a waveform to be open, but with no selection made. Actions begin at the current-time indicator position in the waveform, and they affect any data present at that point.
- Scripts that work on a selection. These scripts require a selection to be made first. Actions in the script apply only to the selection.

A set of scripts can be grouped together in a *script collection*. For example, a script collection called "ambiance" might contain scripts for adding echo, reverb, and delay, and one called "batch utilities" might contain scripts for batch processing. (See ["Batch processing files" on page 251.](#))

Creating scripts

Use the Scripts dialog box to create your scripts.

To create a script:

- 1 Set up Adobe Audition for the script you want to create. For example, open a waveform typical of the ones you'll apply the script to, or, if you want a script that starts from scratch, close all open waveforms.
- 2 In Edit View, choose Options > Scripts. The Script Collections area displays the name of the currently opened script collection. If the collection hasn't been named, the name New Collection appears.
- 3 Do one of the following:
 - To open an existing script collection, click Open/New Collection, navigate to the collection (*.scp) file, and then double-click it.
 - To create a new script collection, click Open/New Collection. Navigate to the folder in which you want to save the new collection (*.scp) file. Then, type a name for it in the File Name text box, and click Open.
 - To rename a script collection, click Edit Script File. The collection (*.scp) file opens in Windows Notepad. Locate the "Collection:" entry on the first line, and type a new name. Then, save the file.

Note: The name in the Script Collections area doesn't reflect the change until you reopen the script.

- 4 Type a name for your script in the Title text box.
- 5 Click Record. The Scripts dialog box closes.
- 6 Perform the actions that you want to be part of the script.

Note: Don't open or save a file as part of the actions for the script, since these actions are specific to a particular file. If you make a mistake, return to the Scripts dialog box, click Stop Current Script, click Clear, and start over.

- 7 After you record the script, choose Options > Scripts.
- 8 Click Stop Current Script.
- 9 Type a description for the script in the text area of the dialog box. The description appears when the script is selected.
Note: You can add or edit a description later by clicking *Edit Script File*.
- 10 Click Add to Collection. The script appears in the list at the left.

Running scripts

After you create a script, you can run it on a file, an entire waveform, or part of a waveform, depending on the script type.



To run a single script on a batch of files, use the Batch Processing command.

To run a script:

- 1 Set up Adobe Audition to match the starting point of the script. For example, if you want to run a script intended for a waveform, open a file and select a waveform. If you want to run a script that starts from scratch, close all open waveforms.
- 2 Switch to Edit View, and choose Options > Scripts. The Script Collections area displays the name of the currently opened script collection. If the collection hasn't been named, the name New Collection appears.
- 3 If the script collection you want isn't open, click Open/New Collection. Navigate to the collection (*.scp) file you want, and double-click it.
- 4 Select the script you want to run from the list.
- 5 Set the following options as desired, and then click Run Script:

Pause At Dialogs Stops the script at each dialog box used in the script, so you can modify the settings at those points. Clicking Cancel in any dialog box stops the script, and clicking OK continues it.

Alert When Complete Displays a notice when the script is finished.

Execute Relative To Cursor When running a Works On Current Wave type of script, performs all script operations relative to the original position of the cursor, as opposed to at the current position.

For example, if a script was recorded with the current-time indicator at 0:10:00, selecting this option applies the script at the current cursor position, plus 10 seconds: If the current cursor position is at 0:05:00, the script would start at 0:15:00.



If you're likely to run a script at the current cursor position, record the script with the cursor at a 0:00:00 position, and select this option when you run it.

Script Type Indicates the type of script selected in the scripts list: Script Starts From Scratch works with all files closed; Script Works On Current Wave works on an entire waveform; and Script Works On Highlighted Section works on the selected part of a waveform.

Editing scripts

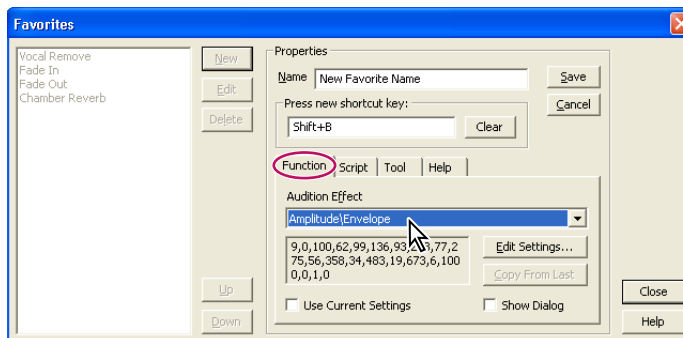
The Edit Script File option in the Scripts dialog box lets you modify existing scripts as text in a Windows Notepad file.

To edit a script:

- 1 In Edit View, choose Options > Scripts. The Script Collections area displays the name of the currently opened script collection. If the collection hasn't been named, the name New Collection appears.
- 2 If the script collection you want isn't open, click Open/New Collection. Navigate to the collection (*.scp) file you want, and double-click it.
- 3 Select the script you want to edit from the list.
- 4 Click Edit Script File. The collection file opens in Windows Notepad.
- 5 Scroll through the file to find the script you want.
- 6 Make the changes you want, and save the file.

Using favorites (Edit View only)

The Favorites menu in Edit View lists custom commands you can create. The Favorites dialog box lets you create, edit, customize, and save these commands, which are based on your favorite Adobe Audition effects, scripts, and even third-party tools (the latter using command line executables). You can even organize favorites into hierarchical submenus for easy navigation.



The four tabs in the Favorites dialog box

To apply favorites:

In Edit View, choose Favorites, followed by the favorite you want to apply.

To create or edit favorites:

1 Choose Favorites > Edit Favorites.

2 Select from the following options, click Save, and then click Close:

New Enables the fields in the Properties area of the Favorites dialog box.

Edit Enables the fields in the Properties area for the selected favorite.

Delete Removes the selected favorite.

Up Moves the selected favorite up in the list. The Favorites menu reflects the order of the list.

Down Moves the selected favorite down in the list.

Name Specifies the name of a favorite. Use this text box to help organize the Favorites menu by doing one or more of the following:

- Create hierarchical menus by using a backslash (“\”). For example, type **My Delays\Hall Reverb** in the Name text box to place the Hall Reverb favorite in the My Delays submenu.
- Add separator bars by typing a series of dashes (“-----”) into the Name text box. If you want more than one separator bar, type a different number of dashes, or add text so that the separator doesn’t match one in the list. For example, type “-----2” (the “2” after the dashes doesn’t appear in the Favorites menu).
- Create a separator bar for a submenu by entering the submenu path first (such as “My Effects\-----”). (The text that appears with a separator bar is for appearance only.)

Note: If you create text for a submenu title, make sure not to specify any command, script, or tool listed on the Function tab, Script tab, or Tool tab.

Press New Shortcut Key Lets you type a key or combination of keys to use as the keyboard shortcut to a favorite. Adobe Audition accepts most single key shortcuts (the most notable exceptions are the Print Screen, Scroll Lock, Number Lock, Caps Lock, Tab, Function, and Enter keys), and it also accepts the Ctrl, Shift, and Alt keys (or any combination of the three) as the first in a combination of keys.

Note: If the keyboard shortcut you type is already used by Adobe Audition, a dialog box appears, giving you the option to overwrite the current shortcut.

Clear Clears text from the Press New Shortcut Key text box.

Function tab Lets you specify the following options:

- Audition Effect lets you choose any command listed in the Effects and Generate menus. After you choose a command, the settings last used for it appear.
- Edit Settings displays the window that corresponds to the command you chose. You can then specify the settings to be used when you choose the favorite from the Favorites menu.
- Copy From Last applies the settings used the last time the particular command was completed successfully.
- Use Current Settings applies the settings currently specified for the particular command. Deselect this option to edit the settings.
- Show Dialog causes the dialog box for the particular command to display, with the settings you specified for the favorite.

Script tab Lets you specify the following options:

- Script Collection File displays the current script collection in use. The button to the right of the text box opens the Browse For Script dialog box that lets you navigate to and select a script collection (*.scp) file.
- Script lets you choose the script you want to run from the selected collection.
- Pause At Dialogs stops the script at each dialog box used in the script, so you can modify the settings at those points. Otherwise, the script runs nonstop to completion.


Tool tab Specifies the command line for the tool you want to run, including any command line switches the particular tool may need. The button to the right of the text box opens the Browse For Tool dialog box that lets you navigate to the desired tool.

Help tab Displays instructions for adding separators and submenus to the Favorites menu.

Burning Audio CDs

Using CD Project View

CD Project View provides an easy-to-use interface for assembling CD tracks, setting track properties, and burning CDs. The display window in CD Project View contains the track list, which displays information about the audio tracks you assemble. CD Project View also shares many elements with Edit View and Multitrack View, such as dockable windows, menus, toolbars, and a status bar. (See [“About the work area” on page 19.](#))



CD Project View

You use the tabs above the display window, commands in the View menu, or buttons in the toolbar to switch between views. (See [“Switching between views” on page 20.](#))

Assembling tracks

You can assemble the tracks for a CD all at once, or you can insert individual tracks as you finish editing the audio. After you insert tracks, you can also change their order or remove them.

When you assemble audio for a CD, you'll probably want to fine tune the individual tracks so that they form a cohesive whole. This process—known as *mastering*—often involves cropping files, adjusting dynamics (compressing), and comparing the audio for continuity levels and EQ. (See [“About the mastering process” on page 101](#) and [“Normalizing groups of files” on page 250.](#))


Inserting tracks

Adobe Audition provides a variety of ways to insert tracks into CD Project View. Keep in mind that you're not limited to inserting entire files—you can also insert audio ranges that are defined as track cues. For more information on creating track cues, see [“Defining and selecting cues” on page 83](#) and [“Choosing a cue type” on page 83.](#)



To insert a track:

Do one of the following:

- Select one or more files or track cues in the Files tab of the Organizer window. Then, drag the items into the track list or click the Insert Into CD Project button . For more information on using the Files tab of the Organizer window, see [“Organizing files” on page 30](#).
- In CD Project View, choose Insert > Audio or Insert > Audio From Video. Select a file, and click Open.
- In CD Project View, choose Insert > File/Cue List. Select the file or track cue you want to insert.
- Drag any supported audio file type from your desktop (Windows, My Computer, or Windows Explorer) directly into the track list in CD Project View. The file first opens in Adobe Audition, and then is inserted into the track list.
- In Edit View, open a file. To insert the entire file, make sure that no audio is selected; to insert part of a file, select the desired range. Then choose Edit > Insert In CD Project.
- In Multitrack View, open a session file, and choose Edit > Mix Down To CD Project. If the session includes track cues, each cue range is inserted into the track list as a separate track.



If you want to divide a single, long audio file (such as a recording of a concert that includes several songs) into multiple tracks on a CD, insert the file into a session, and add track cues at the desired locations. Then, choose Edit > Mix Down To CD Project. The cue ranges are inserted automatically as separate tracks.

Selecting tracks

In CD Project View, you can select one or more tracks by clicking in the track list. You can also select all tracks by choosing the Select All Tracks command.

To select a track:

In CD Project View, click the track in the track list.

To select multiple tracks:

Do one of the following:

- To select adjacent (contiguous) tracks, click the first track in the desired range, and then Shift-click the last.
- To select nonadjacent (incontiguous) tracks, Ctrl-click them.

To select all tracks:

Choose Edit > Select All Tracks.

Rearranging tracks

In CD Project View, you can move tracks up and down to change their play order on a CD.

To rearrange tracks:

- 1 In CD Project View, select the track you want to move.
- 2 Click Move Up or Move Down.

Removing tracks

In CD Project View, you can remove a single track, multiple tracks, or all tracks.

To remove tracks:

In CD Project View, do one of the following:

- Select one or more tracks, and click Remove. Alternatively, choose Edit > Remove Selected Tracks.
- To remove all tracks, click Remove All. Alternatively, choose Edit > Remove All Tracks.

To close the source files when removing tracks:

- 1 Select one or more tracks.
- 2 Choose Edit > Destroy Selected Tracks (Remove and Close).

Editing the source audio for tracks

The Edit Waveform command in CD Project View lets you edit the source audio for a track in Edit View.

To edit the source audio for a track:

- 1 In CD Project View, select the track you want to edit.
- 2 Choose Edit > Edit Waveform.

Setting track properties

Adobe Audition lets you specify a title and artist for each track. CD players that support CD Text display the text during playback.

You can also change the length of pauses between tracks, enable or disable copy protection and pre-emphasis features, and add an ISRC (International Standard Recording Code) number.

To set track properties:

- 1 In CD Project View, select the track for which you want to set properties, and click Track Properties. Alternatively, select the track, and choose View > Track Properties.
- 2 Enter a track title and artist for the track.

Important: In order for Adobe Audition to write text to the CD, you must select Write CD-Text in the Write CD dialog box. (See [“Writing a CD” on page 262.](#))

- 3 If you want to set additional properties for the track, select Use Custom Track Properties. Set any of the following options, and click OK:

Pause Adds a pause of the specified length before the track. By default, Adobe Audition assigns a 2-second pause to the beginning of each track.

Copy Protection Sets the copy protection flag (as defined by the Red Book specification) for the track. In order for copy protection to occur, the CD player must support the copy protection flag.

Pre-Emphasis Sets the pre-emphasis flag (as defined by the Red Book specification) for the track. *Pre-emphasis* is a basic noise reduction process that is implemented by a CD player. For pre-emphasis to occur, the CD player must support the pre-emphasis flag.

ISRC Specifies an ISRC (International Standard Recording Code). This code is used only on CDs that are destined for commercial distribution. ISRC codes have 12 characters and use the following format:

- ISO Country: 2 digit code (for example, *US* for USA).
- Registrant code: 3 digit alpha-numeric, unique reference.
- Year of reference: last 2 digits of the year (for example, *04* for 2004).
- Designation code: a 5 digit, unique number.

Same For All Tracks Applies the settings, with the exception of the ISRC code, to all tracks in the track list.

Writing a CD

Before you write a CD, you should verify that your CD burning device is set up correctly. Then, set CD options and write the CD.

Note: Audio on CDs must be 44.1 kHz, 16 bit, stereo. If you insert a track with a different sample type, Adobe Audition automatically converts the audio for you.

To set CD device properties:

- 1 In CD Project View, choose Options > Device Properties.
- 2 Select the device you want to set up.
- 3 Select a buffer size and write speed for the device.
- 4 If the device supports buffer underrun protection, select Buffer Underrun Protection to allow the drive to stop and resume burning as needed.
- 5 Click OK.

To set CD options and write a CD:

- 1 Insert a blank, writable CD into the CD burning device.
- 2 In CD Project View, click Write CD or choose File > Write CD.
- 3 Choose the device you want to use to write the CD. (Click Device Properties to set device properties, as described in the previous procedure.)
- 4 Choose a setting from the Write Mode pop-up menu:
 - Write CD writes the CD without testing for buffer underruns.
 - Test Write Only tests if the CD can be written without the occurrence of buffer underruns. No audio is written to the CD.
 - Test and Write CD tests for buffer underruns and then proceeds with the actual write process if the test is successful.

5 Select Eject Disc When Complete to eject the CD tray upon completion of the write process.

6 Select Write CD-Text if you want to write text, including the track title and artist for each track, to the CD. Type the desired information in the text boxes for Title, Artist, and UPC/EAN.

Note: *The UPC/EAN is a 13-digit code that is used to uniquely identify merchandise and communicate product information between a vendor and retailer.*

7 Click Write CD. The Track and Disk bars show you the progress of the write process.

Keyboard Shortcuts

About keyboard shortcuts

The default keyboard shortcuts address most audio production needs, but you can also create custom shortcuts tailored to your working style. To customize shortcuts or trigger commands with a MIDI keyboard, use the Keyboard Shortcuts & MIDI Triggers command. (See [“Using shortcuts” on page 21.](#))

Note: Adobe Audition displays most default keyboard shortcuts in menu commands and tool tips. The user guide and Help list only shortcuts that Adobe Audition doesn't display.

Keys for playing audio

Space	Toggle between Play and Stop
Ctrl+Space	Toggle between Record and Pause
Ctrl+Shift+Space	Toggle between Play All and Pause
Alt+O	Play postroll
Alt+R	Play preroll and postroll (skip selection)
Alt+E	Play preroll and selection
Home	Move the current-time indicator to the beginning of the waveform or session
End	Move the current-time indicator to the end of the waveform or session
Page Up	Move the current-time indicator one page to the left
Page Down	Move the current-time indicator one page to the right
Left Arrow	Move the current-time indicator to the left
Right Arrow	Move the current-time indicator to the right



Keys for selecting ranges, channels, and tracks

Up Arrow	Select the left channel or next higher track
Down Arrow	Select the right channel or next lower track
Shift+Home	Extend the selection to the beginning of the waveform or session
Shift+End	Extend the selection to the end of the waveform or session
Shift+Page Up	Extend the selection one page to the left
Shift+Page Down	Extend the selection one page to the right
Shift+Left Arrow	Extend the selection to the left
Shift+Right Arrow	Extend the selection to the right
Ctrl+Shift+A	Select the current page
[Move left side of the selection inward during playback
]	Move right side of the selection inward during playback

Keys for copying waveforms

Ctrl+Insert	Copy the waveform or selection to the clipboard
Shift+Insert	Paste the clipboard's contents into the waveform display or session display
Ctrl+M	Insert the waveform into the session display
Ctrl+Shift+N	Paste the contents of the active clipboard to a new waveform

Keys for editing clips

Ctrl+Up Arrow	Select the previous clip in the currently selected track
Ctrl+Down Arrow	Select the next clip in the currently selected track
Alt+Left Arrow	Nudge the selected clip to the left
Alt+Right Arrow	Nudge the selected clip to the right
Ctrl+Shift+Up Arrow	Clip color (next)
Ctrl+Shift+Down Arrow	Clip color (previous)

Keys for repeating commands

F2	Repeat the last command (its dialog box appears)
F3	Repeat the last command (no dialog box appears)

Keys for using markers

F8	Add a cue or cue range
Shift+F8	Add a CD track marker
Ctrl+F8	Add a CD index marker
1	Mark Intro Time
2	Mark Sec Tone

Keys for scrolling waveforms and sessions

Ctrl+Home	Scroll to the beginning
Ctrl+End	Scroll to the end
Ctrl+Page Up	Scroll one page to the left
Ctrl+Page Down	Scroll one page to the right
Ctrl+Left Arrow	Scroll to the left
Ctrl+Right Arrow	Scroll to the right

Keys for viewing windows


F12	Toggle between Multitrack View and Edit View
Alt+1	Set focus to the main display
Alt+Page Up	Activate the previous floating window
Alt+Page Down	Activate the next floating window
Alt+ /	Flash the window that's currently in focus

Digital Audio Primer

Sound fundamentals

Sound is created by vibrations, such as those produced by a guitar string, vocal cords, or a speaker cone. These vibrations move the air molecules near them, forcing molecules together, and as a result raising the air pressure slightly. The air molecules that are under pressure then push on the air molecules surrounding them, which push on the next set of air molecules, and so forth, causing a wave of high pressure to move through the air. As high pressure areas move through the air, they leave low pressure areas behind them. When these pressure lows and highs—or waves—reach us, they vibrate the receptors in our ears, and we hear the vibrations as sound.

When you see a visual waveform that represents audio, that waveform represents these pressure waves. The zero line in the waveform is the pressure of air at rest. When the line swings up, it represents higher pressure, and when it swings low, it represents lower pressure. This waveform is the equivalent of the pressure waves in the air.



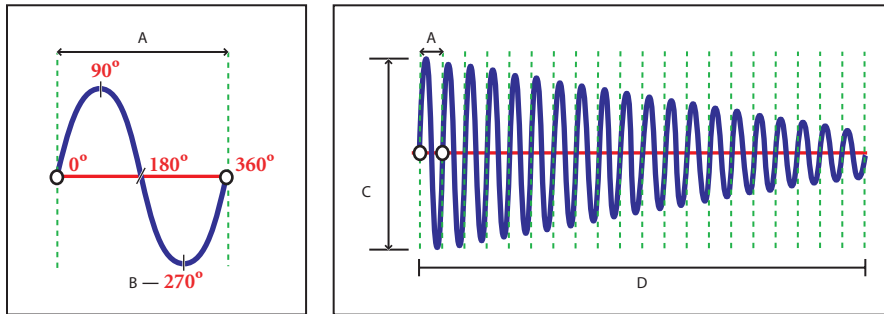
A sound wave represented as a visual waveform

A. Zero line **B.** Low pressure area **C.** High pressure area



Waveforms


Amplitude reflects the change in pressure from the peak of the waveform to the trough. *Cycle* describes the amount of time it takes a waveform to return to the same amplitude level. *Frequency* describes the number of cycles per second, where one Hertz (Hz) equals one cycle per second. That is, a waveform at 1000 Hz goes through 1000 cycles every second. *Phase* measures how far through a cycle a waveform is. There are 360 degrees in a single cycle; if you start measuring at the zero line, a cycle reaches 90 degrees at the peak, 180 degrees when it crosses the zero line, 270 degrees at the trough, and 360 degrees when it completes at zero. *Wavelength* is the distance, measured in units such as inches or centimeters, between two points with the same degree of phase.



A single cycle at left; a 20 Hz waveform at right


A. Wavelength **B.** Degree of phase **C.** Amplitude **D.** One second

When two or more sound waves meet, their amplitudes add to and subtract from each other. If the peaks and troughs of the two waveforms line up, they are said to be in phase. In this case, each peak adds to the peak in the other waveform, and each trough subtracts from the other trough, resulting in a waveform that has higher amplitude than either individual waveform.




In-phase waves reinforce each other.

Sometimes the peaks of one waveform match up with the troughs of another. The peaks and troughs will cancel each other out, resulting in no waveform at all. Such waveforms are said to be 180 degrees out of phase.



Out-of-phase waves cancel each other out.

In all other cases, waves are out of phase by some amount. This results in a waveform that is more complex than either of the original waveforms; continuing to add waves makes a more and more complicated waveform. Keep in mind, however, that a single instrument can create extremely complex waves because of the unique structure of the instrument; a violin and a trumpet sound different even when playing the same note. When you see music, voice, noise, and other complicated sounds represented by a waveform, you see all the waveforms from each sound added together.



Two simple waves combine to create a complex wave.

Analog audio


A microphone works by converting the pressure waves of sound into changes in voltage on a wire. These changes in voltage match the pressure waves of the original sound: high pressure is represented by positive voltage, and low pressure is represented by negative voltage. Voltages travel down the microphone wire and can be recorded onto tape as changes in magnetic strength or onto vinyl records as changes in amplitude in the groove. A speaker works like a microphone in reverse, taking the voltage signals from a microphone or recording and vibrating to re-create the pressure wave.

Digital audio

Unlike analog storage media such as magnetic tape and vinyl records, computers store audio information digitally as a series of zeroes and ones. In digital storage, the original waveform is broken up into individual samples. This process is typically known as *digitizing* or *sampling* the audio, but it is sometimes called *analog-to-digital conversion*. The sampling rate defines how often a sample is taken. For example, CD-quality sound has 44,100 samples for each second of a waveform.

Sampling rate

The sampling rate determines the frequency range of an audio file. The higher the sampling rate, the closer the shape of the digital waveform will be to that of the original analog waveform. Low sampling rates limit the range of frequencies that can be recorded, which can result in a recording that poorly represents the original sound.



Two sample rates

A. Low sample rate that distorts the original sound wave.


B. High sample rate that perfectly reproduces the original sound wave.

To reproduce a given frequency, the sampling rate must be at least twice that frequency. For example, if the audio contains audible frequencies as high as 8000 Hz, you need a sample rate of 16,000 samples per second to represent this audio accurately in digital form. This calculation comes from the Nyquist Theorem, and the highest frequency that can be reproduced by a given sample rate is known as the Nyquist Frequency. CDs have a sample rate of 44,100 samples per second that allows sampling up to 22,050 Hz, which is higher than the limit of human hearing, 20,000 Hz.

Bit depth

Just as the sample rate determines the frequency resolution, the bit depth determines the amplitude resolution. A bit is a computer term meaning a single number that can have a value of either zero or one. A single bit can represent two states, such as on and off. Two bits together can represent four different states: zero/zero, one/zero, zero/one, or one/one. Each additional bit doubles the number of states that can be represented, so a third bit can represent eight states, a fourth 16, and so on.

Amplitude resolution is just as important as frequency resolution. Higher bit depth means greater dynamic range, a lower noise floor, and higher fidelity. When a waveform is sampled, each sample is assigned the amplitude value closest to the original analog wave. With a resolution of two bits, each sample can have one of only four possible amplitude positions. With three-bit resolution, each sample has eight possible amplitude values. CD-quality sound is 16-bit, which means that each sample has 65,536 possible amplitude values. DVD-quality sound is 24-bit, which means that each sample has 16,777,216 possible amplitude values.



Higher bit depths provide greater dynamic range.

Where Adobe Audition fits into the process

When you record audio on your computer, Adobe Audition tells the sound card to start the recording process and specifies what sampling rate and bit depth to use. The sound card determines the supported sample rates and bit depths. Most cards can record and play at CD-quality settings, but many also support other settings (for example, a 48 kHz sample rate, which is common in film and video post-production). Your sound card probably has both Line In and Microphone In ports through which it can accept analog signals. The sound card samples the audio at the specified sample rate and assigns each sample an amplitude value. Adobe Audition stores each sample in sequence until you stop recording. Once you've recorded the audio, you can use Adobe Audition to edit the audio or save it to disk as a file.

When you play a file in Adobe Audition, the process happens in reverse. Adobe Audition tells the sound card that it will play a file and sends a series of digital samples to the card. The sound card reconstructs the original waveform and sends it as an analog signal through the Line Out port to your speakers.

An audio file on your hard drive, such as a WAV file, consists of a small header indicating sample rate and bit depth, and then a long series of numbers, one for each sample. These files can be very large. For example, at 44,100 samples per second and 16 bits per sample, a file includes 705,600 bits per second. This equals 86 kilobytes per second and more than 5 megabytes per minute. Stereo sound has two channels, so CD-quality sound requires a little more than 10 megabytes per minute.

Introducing MIDI

In contrast to a digital audio file, a MIDI file might be as small as 10 kilobytes per minute, so you can store up to one hundred minutes of MIDI per megabyte. MIDI and digital audio are fundamentally different: digital audio is a digital representation of a sound wave, MIDI is a language of instructions for musical instruments. A digital audio file seeks to exactly represent an audio event just like a tape recorder, whether it's a musical performance, a person talking, or any other sound. MIDI, on the other hand, is more like sheet music. It acts as instructions for the re-creation of a musical selection. These MIDI instructions, however, cannot reproduce highly complex sounds, such as the human voice.

MIDI files record information such as the note to be played, the instrument to play the note on, the pan and volume of that particular note, and so on. When a MIDI file is played back, the sound card takes this information and uses its synthesizer to re-create the note on the right instrument. Because every synthesizer sounds different, the MIDI file will sound different when played through different sound cards. MIDI support in Adobe Audition is limited to playback of MIDI files.

Conclusion

To summarize, the process of sampling or digitizing audio starts with a pressure wave in the air. A microphone converts this pressure wave into voltage variations. An analog-to-digital converter, found in devices such as sound cards, samples the signal at the sample rate and bit depth you choose. Once the sound has been transformed into digital information, Adobe Audition can record, edit, process, mix, and save your digital audio files. The possibilities for manipulation of digital audio within Adobe Audition are limited only by your imagination.

Glossary

A

ACM (Audio Compression Manager) A Microsoft technology that enables Windows applications to compress and decompress files in a variety of formats, such as DSP Group TrueSpeech and GSM 6.10. Some ACM formats install with Windows, while others install with software applications.

ADPCM (Adaptive Differential Pulse Code Modulation) An audio compression scheme that compresses sound files from 16 bits to 4 bits, yielding a 4:1 compression ratio. There are many varieties of ADPCM, such as the IMA (Interactive Multimedia Association) DVI standard, and versions from Microsoft, Dialogic, and others.

ActiveMovie See [“DirectX” on page 276](#).

Adapter A cable, plug, or jack that enables you to connect two audio or video devices together.

ADAT A digital 8-track tape deck manufactured by Alesis Corporation that is very popular in recording studios.

Aliasing Noise that occurs when a high frequency sound exceeds the Nyquist Frequency for a given sample rate. (See [“Nyquist Frequency” on page 280](#).) Most analog-to-digital converters prevent aliasing by filtering out sounds above the Nyquist Frequency.

Amplitude Amplitude represents the loudness of an audio signal. A waveform’s amplitude is measured by its distance from the center line, which represents an amplitude of 0. There are different standards for measuring amplitude, but the decibel (dB) is the most common. (See [“Decibel \(dB\)” on page 276](#).)

Analog recording Traditional audio recording with devices such as magnetic tape machines and vinyl records. Analog audio recording consists of a continuous curve, as opposed to digital recording, which consists of discrete samples.

ASCII text data You can represent audio data in this standard text format (.txt), with each sample separated by a carriage return, and channels separated by a tab character. Before the audio data, you can add a header with a format of `Keyword:Value`, with the keywords `Samples`, `BitsPerSample`, `Channels`, `SampleRate`, and `Normalized`. (The values for `Normalized` are `True` or `False`.) If no header exists, the data is assumed to be 16-bit signed decimal integers.

Attack The first part of the sound that you hear. Some sounds (like pianos and drums) have a very fast attack; the loudest portion of the sound occurs very quickly. A sound with a slow attack rate (such as a soft string section) slowly increases in volume.

Attenuate To reduce volume or signal level.



Audio file format The method used to store audio data on disk, chosen in Save dialog boxes. Adobe Audition supports many file formats, and each supports a variety of properties such as sample rate and compression. Some file formats may not be compatible with other platforms. On the Windows platform, Windows PCM (.wav) is the most common format.

Audition loop See [“Audition Loop \(.cel\)” on page 242](#).

Automation The process of recording volume and pan changes during a mix, and perfectly reproducing those changes every time a mix plays. In hardware mixers that support automation, volume and pan controls record timing information and physically move during playback. In Adobe Audition, you automate mixes with visual envelopes. (See [“Envelopes” on page 277](#).)

B

Background mixing The process that Adobe Audition uses to mix audio for playback in Multitrack View. Background mixing occurs behind the scenes, reflecting changes to a session, such as a moved or deleted clip, a volume change, or a newly recorded track. The progress of background mixing is displayed by the Mix Gauge. (See [“Mix Gauge” on page 279](#).)

Band pass filter A filter that allows some audio frequencies to pass through unchanged.

Basic cue One of four types of Adobe Audition cues. Basic cues mark important sections of a waveform for later reference (for example, to identify an editing point). These cues also specify stop and start positions for the Play List. (See [“Play List” on page 280](#).)

Beat cue One of four types of Adobe Audition cues. Beat cues function like basic cues, but they specifically identify musical beats.

Beats per minute (bpm) Musical tempo, which is defined by the number of beats that occur every 60 seconds.

Bit Part of the numbering system used in digital data. Bits are combined in groups to form digital words, which represent the changing amplitude values of an analog signal. Bit resolution describes the number of bits used in each word, determining the number of possible amplitude values. Therefore, higher bit resolutions produce higher dynamic range.

Bit resolution (or bit depth) The number of bits used to represent audio amplitude. 8-bit resolution provides a maximum of 256 unique amplitude levels, producing a 48 dB dynamic range; 16-bit resolution provides 65,536 unique amplitude levels, producing a 96 dB dynamic range. Compact disc players have 16-bit resolution, but some sound cards support resolutions higher than 16-bit. Adobe Audition supports up to 32-bit resolution. For the best audio quality, remain at the 32-bit level while transforming audio in Adobe Audition, and then convert to a lower resolution for output.

Brown noise Brown noise has a spectral frequency of $1/f^2$, so it emphasizes low-frequency components, resulting in thunder- and waterfall-like sounds. Brown noise follows a Brownian motion curve, in which each sample in a waveform contains a mixture of predefined and random frequency components.

Bus In hardware mixers, a channel that lets you combine several other channels and output them together. In Adobe Audition’s Multitrack View, you can similarly use software buses to combine several tracks.

Burn To write to a CD-R or CD-RW disc.

C

CD-R A recordable compact disc that you can write to only once. These discs typically hold 650 MB of data, which equals 74 minutes of stereo audio. *CD-R* sometimes refers to the computer drives that burn CD-R discs.

CD-RW A rewritable compact disc. These discs typically hold 650 MB of data, which equals 74 minutes of stereo audio. Unlike a CD-R, however, a CD-RW disc can be erased and written to again.

Chorus A delay effect that simulates several voices by adding multiple short delays with a medium amount of depth and a small amount of feedback.

Click track An audio track comprised of clicks that occur on the beat, like a metronome. Click tracks are often used at the beginning of a session to provide timing information for musicians and then removed from the session before mixing down.

Clip A visual representation of individual audio, video, or MIDI files in Adobe Audition's Multitrack View.

Clipping In digital audio, distortion that occurs when the amplitude of a signal exceeds the maximum level for the current bit resolution (for example, 256 in 8-bit audio). Visually, clipped audio produces broad flat areas at the top of a waveform. If you experience clipping, lower the recording input or the source output levels.

CODEC (Compressor/Decompressor) An abbreviation often used to describe multimedia compression schemes used by ACM, MPEG, QuickTime, AVI, and the combined A-D-D-A modules on some sound cards.

Compressor Reduces dynamic range by lowering amplitude when an audio signal rises above a specified threshold. For example, compressors can be used to eliminate variations in the level of an electric bass, providing an even, solid bass line. Compressors can also compensate for variations in level produced by a vocalist who moves frequently or has an erratic volume.

Crossfade A fade from one audio track to another.

Crosstalk Undesired leakage of audio from one track to another, a common problem with analog tape. Crosstalk is impossible in Adobe Audition because each track is stored as a separate digital audio file.

Cue List A list of time locations defined in an audio file. A cue can be either a point that specifies a time position or a range that specifies a selection. In Adobe Audition, you can define and save an unlimited number of cues for later recall or for assembly in the Play List window. (See ["Play List" on page 280.](#))

D

DAC (Digital-to-Analog Converter) The hardware responsible for converting a digital audio or video signal into an analog signal that you can play through amplifiers and speakers.

DAT (Digital Audio Tape) A standard two-track digital audio tape format. DAT tapes are sampled at 16 and 24 bits, and 32,000, 44,100, and 48,000 samples per second. (The latter is often described as *DAT quality*).

DC offset Some sound cards record with a slight DC offset, in which direct current is introduced into the signal, causing the center of the waveform to be offset from the zero point (the center line in the waveform display). DC offset can cause a click or pop at the beginning and end of a file. To compensate for DC Offset, use the DC Bias Adjust setting provided by the Amplify command.

Decibel (dB) In audio, the decibel (dB) is a logarithmic unit of measurement used for amplitude.

dBFS Decibels below full scale in digital audio. 0 dBFS is the maximum possible amplitude value (for example, 256 for 8-bit audio). A given dBFS value does not directly correspond to the original sound pressure level measured in acoustic dB.

Delay A time-shifted signal that you can mix with the original, nondelayed signal to provide a fuller sound or create echo effects. Adobe Audition offers a variety of delay effects such as Reverb, Chorus, and Echo.

Destructive editing Editing (such as cutting and pasting, or effects processing) that changes the original audio data. For example, in destructive editing, a change in audio volume alters the amplitude of the original wave file. In Adobe Audition, Edit View is a destructive editing environment; however, edits do not permanently change audio until you save a file.

Devices Wave and MIDI devices that send data into and out of the computer. In Adobe Audition, wave devices are sound card inputs and outputs used for recording and playback of audio; MIDI devices are hardware interfaces used to send performance and synchronization information to Adobe Audition and other MIDI-enabled programs and hardware. You can configure both device types in the Device Properties dialog box.

Digital Signal Processing (DSP) The process of transforming a digital audio signal by using complex algorithms. Examples include filtering with equalizers, and effects processing with reverbs and delays.

DirectX A development platform designed by Microsoft that provides an open standard for audio plug-ins. Plug-ins based on this standard can be used by any application that supports DirectX, such as Adobe Audition.

Dither Dithering adds small amounts of noise to a digital signal so that very quiet audio remains audible when you convert from a high bit resolution to a lower one (for example, when converting from 32-bit to 16-bit). Without dithering, quiet audio passages such as long reverb tails may be abruptly truncated.

Dry Used to describe an audio signal without any signal processing such as reverb; the opposite of Wet.

DSP See [“Digital Signal Processing \(DSP\)” on page 276](#).

DVD (Digital Video Disc) A storage medium similar to a compact disc (CD), but with much higher bandwidth and storage capabilities. Audio stored in DVD movies is generally 96 kHz/24-bit.

E

Echo A distinct repetition of a sound, caused by the sound reflecting off a surface. Adobe Audition offers two echo effects, Echo and Echo Chamber.

8-bit Signed See [“8-bit signed \(.sam\)” on page 240](#).

Envelopes To automate mixes in Multitrack View, Adobe Audition uses envelopes, which are drawn directly on clips. Envelopes visually indicate the pan, volume, wet/dry, and effects parameter settings at any point in a track. For example, when a volume envelope is at the top of an audio clip, the audio is at full volume; when the envelope is at the bottom, the audio is at zero volume.

Equalization (EQ) The process of increasing or decreasing the amplitude of specific audio frequencies relative to the amplitude of other audio frequencies.

Expander Increases dynamic range by lowering amplitude when an audio signal falls below a specified threshold (the opposite of a compressor). For example, an expander can be used to lower the level of background noise that becomes audible when a musician stops playing.

F

Fast Fourier Transform (FFT) An algorithm based on Fourier Theory that Adobe Audition uses for filtering, Spectral View, and Frequency Analysis features. Fourier Theory states that any waveform consists of an infinite sum of sin and cos functions, allowing frequency and amplitude to be quickly analyzed. Higher FFT sizes create more precise results but take longer to process.

Flange An audio effect caused by mixing a varying, short delay in roughly equal proportion to the original signal.

Flushing The process Adobe Audition performs when it copies the audio data from a waveform file to Adobe Audition’s temp folder so that the original file can be closed. This allows the file to be renamed, deleted, or opened exclusively by another application. Flushing sometimes occurs when a modified waveform is saved on top of its original file.

Frequency Measured in Hertz (Hz), cycles per second, frequency describes the rate at which a sound wave vibrates. A cycle consists of movement from a starting point (0) through both positive and negative amplitudes, eventually returning to the starting point. A sound’s frequency determines its pitch: high frequency equals high pitch, and low frequency equals low pitch.

FX An abbreviation for *effects*.

H

Hertz (Hz) Cycles per second. A unit of measurement that describes the frequency of a sound. (See [“Frequency” on page 277](#).)

I

Index cue One of four types of Adobe Audition cues. Index cues become index markers in a CD track. If a CD player is configured to display remaining time, it displays the time before track markers and index markers. Note, however, that not all CD players support index markers.

Impulse A data file that the Convolution effect uses to modify samples. Impulses function like amplitude maps. For example, if you apply an impulse of a single full-volume sample, the original audio data will be unchanged. Should the impulse be at half volume, however, the original audio data will be reduced to half volume. If several such impulses occur over time, each with descending amplitude, the original audio data will gradually and rhythmically become lower in volume.

Interpolate To estimate the values of data points between known data points. Interpolation is used when new data must be generated to fill in areas where values are unknown.

L

Level Meters Adobe Audition's Level Meters are found by default along the bottom of the application window, and they are used to monitor the volume of incoming and outgoing signals. The red clip indicator to the right of the meters will light up and remain lit when levels exceed the maximum of 0 dB. Clicking the clipping indicator resets it. The top meter represents the left channel, and the bottom meter represents the right.

Limiter A signal processor that limits input signals that exceed a specified threshold level. Above the threshold, the output level remains constant even if the input increases in volume.

Loop An audio file that contains tempo and pitch information, allowing it to match the tempo and pitch of other loops in a multitrack session. You can repeat a loop-enabled clip infinitely by simply dragging its bottom right corner.

M

Mastering The process of finalizing audio for a specific medium, such as the Web or audio CD. Mastering consists of several processing phases, with equalization and compression phases being the most essential. You can master audio files either individually or in groups. (Collectively mastering groups of files is particularly important if the destination medium is audio CD.)

MIDI Musical Instrument Digital Interface, a way of communicating performance information from one piece of software or hardware to another. MIDI can simply relay musical notes, or it can transmit detailed information about timing, synthesizer patches, and such. Windows transmits MIDI information internally between applications; to transmit MIDI information to and from your computer and external devices such as MIDI keyboards, you must use a hardware MIDI interface (for example, the MIDI In port of a sound card).

MIDI Timecode (MTC) A method of sending timing information between MIDI-capable devices. For example, you can convert SMPTE timecode to MTC to synchronize Adobe Audition's transport controls with a video or audio tape deck.

MIDI Trigger An Adobe Audition shortcut triggered by a MIDI event, such as Note On. You can send MIDI events to any device capable of issuing a MIDI command, such as MIDI keyboards and sequencers.

Millisecond (ms) One thousandth of a second. (There are 1000 milliseconds in a second.)

Miniplug A common name for 1/8-inch plugs and jacks, sometimes known as *minijacks*. On the most common sound cards, miniplug jacks provide analog audio inputs and outputs.

Mixdown The process of combining the output of several tracks in Multitrack View to create a new stereo waveform. When you mix down, track properties such as Volume and Pan are reflected in the resulting waveform, so mixdown is typically performed when you're happy with the sound of a session. A mixdown can also produce submixes of selected tracks. For example, you could create a submix of multiple drum tracks and place it on a single, open track, cleaning up the Multitrack View workspace.

Mixing The process of combining multiple audio sources or tracks together for output as a single source. Output is generally in the form of a stereo pair of channels, though mixes may be directed to any number of channels for output (for example, one channel for monophonic output, or 6 channels for surround-sound output).

Mix Gauge Found below the track controls area in Multitrack View, the Mix Gauge indicates the progress of background mixing. Whenever you edit a session, the Mix Gauge becomes blank and then gradually fills as the mix is reprocessed, turning brighter in color when background mixing is complete. You don't need to wait for the Mix Gauge to finish before playing a session.

Mono A monophonic signal, which contains only one sound source.

N

Noise gate A special type of expander that reduces or eliminates noise by greatly lowering signal levels that fall below a specified threshold. Noise gates are often configured to totally eliminate background noise during musical pauses. You can also use these gates to silence pauses in speech.

Noise shaping A technique that shifts the frequency of dithering noise to minimize its audibility.

Nondestructive editing Nondestructive edits don't alter a sound file on disk in any way. For example, nondestructive volume changes do not alter the amplitude of a waveform, but instead simply instruct an audio application to play the waveform at higher volume. In Adobe Audition, Multitrack View is a nondestructive editing environment.

Normalize To adjust the highest peak of a waveform to a certain percentage relative to the digital maximum, 0 dBFS, thereby raising or lowering all other peaks accordingly. Typically, audio is normalized to 100% to achieve maximum volume, but Adobe Audition lets you normalize to any percentage.

Nyquist Frequency Also called Nyquist Rate, this frequency equals half the current sample rate and determines the highest reproducible audio frequency for that sample rate. For example, audio CDs use a sampling rate of 44,100 Hz because the resulting Nyquist Frequency is 22,050 Hz— just above the limit of human hearing, 20,000 Hz. Likewise, to reproduce a signal with an 11,000 Hz frequency range, you must use a sample rate of at least 22,000 Hz. To avoid aliasing distortion, nearly all analog-to-digital converters filter out frequencies that exceed the Nyquist Frequency before the analog-to-digital conversion process. For the best audio quality, record and edit at higher sample rates and then convert down if needed.

O

Offline editing See [“Destructive editing” on page 276](#).

Order A value that determines the slope of an audio filter. First-order filters attenuate an additional 6 dB per octave, second-order filters attenuate 12 dB, third-order filters 18 dB, and so on.

P

PCM (Pulse Code Modulation) PCM is the standard method used to digitally encode audio and is the basic, uncompressed data format used in file formats such as WAV and AIFF.

Peak files Cache files with the extension .pk that enable Adobe Audition to open, save, and redraw audio files more quickly. You can safely delete peak files or deselect the Save Peak Cache Files option in the Settings dialog box. However, keep in mind that without peak files, larger audio files will reopen more slowly.

Phase The position of a sound wave relative to other sound waves. As a sound wave travels through the air, it compresses and expands air molecules in peaks and troughs, much like an ocean wave. In the waveform display, peaks appear above the center line, troughs appear below. If two channels of a stereo waveform are exactly opposite in phase, they will cancel each other out. More common, however, are slightly out-of-phase waves, which have misaligned peaks and troughs, resulting in duller sound.

Pink noise Noise with a spectral frequency of $1/f$, producing the most natural-sounding generated noise. By equalizing pink noise, you can simulate rainfall, waterfalls, wind, a rushing river, and other natural sounds. On the audio spectrum, pink noise falls exactly between brown and white noise.

Play List An arrangement of Cue List entries that you can play in any order and loop a specified number of times in nondestructive fashion. Adobe Audition saves Play Lists in the header of WAV files.

Plug-in A software component that you can add to another piece of software to increase its functionality. Adobe Audition supports third-party VST and DirectX audio plug-ins, which seamlessly integrate into Adobe Audition’s interface.

Preset Most dialog boxes in Adobe Audition support presets, which are settings saved under a particular name for later recall. Dialog boxes that support presets have a Preset list where you can click a preset to recall its settings, and Add and Del buttons for creating and deleting presets.

Preview Many dialog boxes in Adobe Audition offer real-time Preview buttons, letting you monitor setting changes as you make them. The preview quality depends upon your system's performance.

Punch in A recording method used to insert a new recording into a specific region of an existing waveform, usually to replace an undesirable section. Adobe Audition supports punch-in recording in Multitrack View and allows for multiple takes; you can repeatedly record over the original material and afterward choose the best performance.

Q

Quantization A process that occurs when an analog waveform is converted to digital data and becomes a series of samples. Quantization noise is introduced as some samples are shifted to quantization levels allowed by the current bit resolution. This noise is highest at low bit resolutions, where it can particularly affect low amplitude sounds.

R

RCA cable Sometimes called a phono cable, RCA cables have RCA plugs or jacks at either end and are normally used to connect stereo system components, such as receivers, CD players, and cassette decks.

Real time In computer-based audio, *real time* refers to functions that react immediately to user input and transform audio nondestructively. (Note, however, that system speed ultimately determines processing time.) Adobe Audition provides real-time mixing and effects in Multitrack View, and real-time effects previews in Edit View.

Referenced clip In Multitrack View, a referenced clip shares a source file with other clips. For example, if a drum hit occurs 30 times in a session, you can conserve disk space by using 30 referenced clips of the same source file. Because referenced clips represent the same file, any alteration to a referenced clip (like a cut or transform) affects all instances in a session. By contrast, unique clip copies create a separate sound file on disk, consuming more disk space, but allowing for separate editing.

Resample To convert a sound file to a different sample rate.

Reverb The reverberant sound produced by an acoustic space, such as a room or concert hall. Reverb consists of dense, discrete echoes that arrive at the ear so rapidly that the ear can't separate them. Adobe Audition offers four reverb effects: Quick Verb, Studio Reverb, Reverb, and Full Reverb.

Rip The process of digitally extracting audio from a compact disc and turning it into a waveform. Most newer CD-ROM, CD-R, and CD-RW drives support digital audio extraction.

RMS (Root-mean-square) A mathematical formula used to determine the average amplitude of an audio selection. RMS amplitude reflects perceived loudness better than peak amplitude.

S

S/N ratio Signal-to-noise ratio describes the difference between the highest signal level before distortion and the average level of the noise floor. In most analog systems, such as microphone preamps, the S/N ratio is around 92 dB.

Sample A digital snapshot of an audio waveform at a particular point in time. In digital audio, a series of numeric samples reproduces an entire waveform, with higher sample rates producing increased frequency response. (Note that musical samplers use the term *sample* to describe a digital recording, rather than a digital snapshot.)

Sample rate The number of samples per second. Higher sample rates produce increased frequency response but require more disk space. To reproduce a given audio frequency, the sample rate must be at least double that frequency. (See [“Nyquist Frequency” on page 280.](#))

Sampler A musical device that records and plays digital sounds (known as *samples* in this context) and lets you edit and store those sounds.

Sequencer A programmable electronic device that can record and play a sequence of musical events, such as samples, pitches, and rests. Most modern sequencers are MIDI-based. (See [“MIDI” on page 278.](#))

Session A multitrack project in Adobe Audition. Session files are stored with the extension .ses and contain details such as mixing and effects settings. Session files don't contain audio data; instead they contain pathnames pointing to the sound files used in the session.

64-bit Doubles See [“64-bit doubles \(RAW\) \(.dbl\)” on page 240.](#)

SMPTE timecode (Society of Motion Picture and Television Engineers timecode) A timing reference used to synchronize two devices. SMPTE timecode is divided into hours, minutes, seconds, and frames.

Sound card A hardware device that lets your computer play and record audio.

Sound wave A wave of air molecules. Humans can hear sound waves with frequencies of 20 to 20,000 Hz.

Stereo A signal with a left and right channel, allowing for spatial placement of sounds.

Stripe To copy SMPTE timecode to a single track of a multitrack tape so remaining tracks can be synchronized with other devices.

T

Tempo The rhythmic speed of music, normally measured in bpm. (See [“Beats per minute \(bpm\)” on page 274.](#))

Timecode An audio or digital signal that synchronizes time between multiple devices. The most common forms are SMPTE and MIDI timecode.

Track A container for one or more clips in Multitrack View. Each track has independent settings for volume, pan, EQ, effects, and input and output. Each session can have up to 128 tracks.

Track cue One of four types of Adobe Audition cues. Track cues indicate start points for CD tracks.

Track controls The area of Multitrack View that controls each track, with independent settings for volume, pan, EQ, effects, and input and output.

TXT See [“ASCII text data” on page 273](#).

U

Unity gain An amplification level that precisely corresponds to the input signal level, without amplifying or lowering it. (Note that audio hardware operates at two line levels: –10 dBV for consumer equipment, and +4 dBu for professional. If these two hardware types are connected, unity gain will result in a lowered input for consumer equipment, and a raised input for professional.)

W

Wave file Any audio file format that contains primarily sound wave data. Wave files can be in formats such as WAV, AU, AIF, or mp3.

Waveform A term that describes the visual representation of an audio signal, displayed as amplitude across time in Adobe Audition. (In acoustics, *waveform* refers to a sound wave of a specific frequency.)

Waveform clip A visual representation of a wave file or related image in Multitrack View. Edits of these clips are nondestructive.

Waveform display The area of Edit View in which you view and edit audio data. By default this audio material appears as a waveform, but you can view it in spectral form by choosing View > Spectral View.

Wet Used to describe an audio signal that includes signal processing such as reverb; the opposite of Dry.

White noise White noise has a spectral frequency of 1, so equal proportions of all frequencies are present. Because more individual frequencies exist in the upper ranges of human hearing, white noise sounds very hissy. Adobe Audition generates white noise by choosing random values for each sample.

Z

Zero crossing A point in time where a waveform crosses the zero amplitude line. To make edits sound smoother, place them at zero-crossing points, thus avoiding abrupt changes in amplitude that cause pops and clicks.

Legal Notices

Copyright

© 2004 Adobe Systems Incorporated. All rights reserved.

Adobe® Audition™ 1.5 User Guide for Windows®.

If this guide is distributed with software that includes an end-user agreement, this guide, as well as the software described in it, is furnished under license and may be used or copied only in accordance with the terms of such license. Except as permitted by any such license, no part of this guide may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, recording, or otherwise, without the prior written permission of Adobe Systems Incorporated. Please note that the content in this guide is protected under copyright law even if it is not distributed with software that includes an end-user license agreement.

The content of this guide is furnished for informational use only, is subject to change without notice, and should not be construed as a commitment by Adobe Systems Incorporated. Adobe Systems Incorporated assumes no responsibility or liability for any errors or inaccuracies that may appear in the informational content contained in this guide.

Please remember that existing artwork or images that you may want to include in your project may be protected under copyright law. The unauthorized incorporation of such material into your new work could be a violation of the rights of the copyright owner. Please be sure to obtain any permission required from the copyright owner.

Any references to company names in sample templates are for demonstration purposes only and are not intended to refer to any actual organization.

Adobe, the Adobe logo, Adobe Audition, Adobe Encore DVD, Adobe Premiere, and After Effects are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries. Apple, Macintosh, and Mac OS are trademarks of Apple Computer, Inc., registered in the U.S. and other countries. Microsoft, Windows, and Windows NT are registered trademarks of Microsoft Corporation in the U.S. and/or other countries. mp3PRO audio coding technology licensed from Coding Technologies, Fraunhofer IIS and Thomson Multimedia. VST is a trademark of Steinberg Media Technologies AG. ReWire is a product of Propellerhead Software. All other trademarks are the property of their respective owners.

Supply of this product does not convey a license nor imply any right to distribute MP3-encoded or mp3PRO-encoded data created with this product in revenue-generating broadcast systems (terrestrial, satellite, cable, and/or other distribution channels), streaming applications (via Internet, intranets, and/or other networks), other content distribution systems (pay-audio or audio-on-demand applications and the like) or on physical media (compact discs, digital versatile discs, semiconductor chips, hard drives, memory cards, and the like). An independent license for such use is required. For details, please visit <http://mp3licensing.com>

Notice to U.S. government end users. The software and documentation are "Commercial Items," as that term is defined at 48 C.F.R. §2.101, consisting of "Commercial Computer Software" and "Commercial Computer Software Documentation," as such terms are used in 48 C.F.R. §12.212 or 48 C.F.R. §227.7202, as applicable. Consistent with 48 C.F.R. §12.212 or 48 C.F.R. §§227.7202-1 through 227.7202-4, as applicable, the Commercial Computer Software and Commercial Computer Software Documentation are being licensed to U.S. Government end users (a) only as Commercial Items and (b) with only those rights as are granted to all other end users pursuant to the terms and conditions herein. Unpublished-rights reserved under the copyright laws of the United States. Adobe Systems Incorporated, 345 Park Avenue, San Jose, CA 95110-2704, USA. For U.S. Government End Users, Adobe agrees to comply with all applicable equal opportunity laws including, if appropriate, the provisions of Executive Order 11246, as amended, Section 402 of the Vietnam Era Veterans Readjustment Assistance Act of 1974 (38 USC 4212), and Section 503 of the Rehabilitation Act of 1973, as amended, and the regulations at 41 C.F.R. Parts 60-1 through 60-60, 60-250, and 60-741. The affirmative action clause and regulations contained in the preceding sentence shall be incorporated by reference.



Index

A

absorptive surfaces, simulating with Echo effect [158](#)

ACM

defined [273](#)

Waveform format [241](#)

adapter, defined [273](#)

ADAT, defined [273](#)

Add/Remove VST Directory command [36](#)

Adjust Boundaries command [189](#)

Adjust Sample Rate command [94](#)

Adobe After Effects, remixing soundtracks from [222](#)

Adobe Audition

and basic audio workflow [271](#)

tutorials [9](#)

Adobe Premiere Pro, remixing soundtracks from [222](#)

ADPCM

compression [243](#)

defined [273](#)

Advanced Session Properties dialog box [184](#)

AIFF format [242](#)

A-Law Wave format [241](#)

aliasing, defined [273](#)

Allpass Feedback option [162](#)

Amiga IFF-8SVX format [241](#)

amp overdrive [172](#)

Amplify/Fade effect [130](#)

amplitude

defined [273](#)

optimizing [130](#)

analog audio, fundamentals of [269](#)

analog recording, defined [273](#)

Analyze Loudness tab, Group Waveform Normalize dialog box [250](#)

Append To Session command [58](#)

Apple AIFF format [242](#)

ASCII text data

about [242](#)

defined [273](#)

ASPI (Advanced SCSI Programming Interface) [45](#)

attack, defined [273](#)

attenuate, defined [273](#)

AU format [246](#)

audio

analyzing [102](#)

copying [80](#)

mastering [101](#)

muting [87](#)

restoring [101](#)

shrinking or stretching [154](#), [195](#)

Audio Clip Properties window [193](#)

audio digital extraction [59](#)

audio file format, defined [274](#)

audio files

See also specific format names

exporting mixdowns to [239](#)

inserting in sessions [58](#)

saving [237](#)

audio tracks, about [197](#)

Audition Loop format [242](#)

Auto Click/Pop Eliminator effect [108](#)

Auto-Cue feature [85](#)

automation, defined [274](#)

auto-play [56](#), [69](#)

AVI format, exporting mixdowns to [239](#)

B

background mixing

defined [274](#)

priority level [51](#)



band pass filter, defined [274](#)
Bars And Beats command [220](#)
Bars and Beats time format
 about [64](#)
 tempo and [216](#)
Basic cue
 about [83](#)
 defined [274](#)
Bass Management circuit [231](#)
batch processing
 about [249](#)
 cue list [249](#)
 cue ranges [249](#)
 normalizing files [250](#)
 with scripts [251](#)
Batch Processing command [252](#)
beat cue, defined [274](#)
beats
 finding [78](#)
 marking with cues [83](#)
beats per minute (bpm)
 and tempo [216](#)
 defined [274](#)
Binaural Auto-Panner effect [138](#)
binaural cues, adding with Chorus effect [147](#)
bit depth
 automatically converting files to 32-bit [49](#)
 changing [96](#)
 defined [274](#)
 downsampling for playback [39](#)
 fundamentals of [270](#)
bit resolution, defined [274](#)
bit, defined [274](#)
brown noise
 about [91](#)
 defined [274](#)
buffers
 playback, Edit View [45](#)
 playback, Multitrack View [51](#)
 real-time preview [45](#)
 recording, Edit View [45](#)
 recording, Multitrack View [51](#)

burn, defined [274](#)
burning CDs
 in Adobe Audition [262](#)
 tips when using another burning program
 [237](#), [239](#)
Bus Mixer [211](#)
Bus Properties dialog box [211](#)
bus, defined [274](#)

C

Cart Chunk information [100](#)
CBR [235](#)
CBR encoding [244](#), [247](#)
CCITT formats [241](#), [246](#)
CD Project View
 switching to [20](#)
 using [259](#)
CD-R, defined [275](#)
CD-RW, defined [275](#)
CDs
 adding text [261](#), [263](#)
 editing source audio [261](#)
 extracting tracks from [59](#)
 inserting tracks [259](#)
 pre-emphasis [262](#)
 properties for CD-burning devices [262](#)
 properties for tracks [261](#)
 recording from [61](#)
 removing tracks [261](#)
 selecting tracks [260](#)
 setting pauses between tracks [261](#)
 writing [262](#)
CEL format [242](#)
Center Channel Extractor effect [142](#)
center lines, showing or hiding [48](#)
Channel Mixer effect [139](#)
channels
 converting between mono and stereo [95](#)
 selecting data in left or right [76](#)
 specifying which to edit [79](#)
Check for Hidden Clips command [196](#)
Chorus effect [145](#)

Click/Pop Eliminator effect [109](#)
clip indicators, clearing [71](#)
Clip Restoration effect [113](#)
clipboard
 choosing [79](#)
 clearing on exit [46](#)
 resampling data [50](#)
clipping
 displayed in Level Meters [70](#)
 showing or hiding clipping boundaries [49](#)
 when normalizing groups of files [251](#)
clips
 about [187](#)
 aligning [189](#)
 changing properties [193](#)
 copying [191](#)
 crossfading [194](#)
 editing [189](#)
 envelopes for [207](#)
 grouping [188](#)
 inserting in sessions [58](#)
 locking in time [193](#)
 mixing down [213](#)
 moving [188](#)
 muting [193](#)
 protecting from recording [193](#)
 rejoining [191](#)
 removing from a session [196](#)
 repeating [192](#)
 selecting [188](#)
 snapping [189](#)
 splitting [191](#)
 time stretching [195](#)
 video [223](#)
closing files [30](#), [240](#)
Coloration tab, Full Reverb effect [165](#)
command line, playing audio from [69](#)
commands, choosing [21](#)
consistent volume [250](#)
Constant Bit Rate (CBR) [235](#)
context-sensitive menus [21](#)
control signal, for Vocoder effect [177](#)

Convert Sample Type command
 using to change the bit depth [96](#)
 using to change the sample rate [94](#)
 using to convert between stereo and mono [95](#)
 using to convert sample rate of multiple files [97](#)
converting
 audio formats [252](#)
 sample rate of session [238](#)
 sample type [56](#)
 to a different format [237](#), [238](#)
Convolution effect [169](#)
copy protection for CDs [261](#)
Copy Reference Here command [191](#)
Copy To New command [215](#)
Copy Unique Here command [191](#)
copying audio [80](#)
Creative Labs Audigy [233](#)
Creative Sound Blaster format [243](#)
Crossfade commands [194](#)
crossfades, duration of [52](#)
crossover frequency points [236](#)
Cue List
 batch processing [249](#)
 defined [275](#)
Cue List command [249](#)
cues
 adding silence between [85](#)
 adjusting [84](#)
 creating [83](#)
 deleting [85](#)
 inserting in a session [58](#)
 merging [84](#)
 playing [83](#)
 point cues versus range cues [82](#)
 renaming [84](#)
 saving [237](#), [239](#)
 saving to files [85](#)
 selecting [83](#)
 setting automatically [85](#)
 showing or hiding cue lines [48](#)
 snapping to [79](#)
 types of [83](#)

current-time indicator, about [62](#)
custom time format [64](#)
custom timecode display [44](#)
cutting audio [80](#)

D

damping
 and Echo Chamber effect [160](#)
 and Studio Reverb effect [169](#)
DC bias
 adjusting with Amplify/Fade effect [131](#)
 adjusting with Normalize effect [137](#)
 removing when recording [40](#)
DC offset
 adjusting [71](#)
 viewing [108](#)
decay
 adjusting for Echo effect [158](#)
 adjusting for QuickVerb effect [166](#)
 adjusting for Studio Reverb effect [169](#)
Default Session commands [181](#)
Delay effect [156](#)
delay units, and Multitap Delay effect [161](#)
delay, adding to audio [156](#)
Delete Silence command [88](#)
Delete This Take command [67](#)
deleting
 audio [82](#)
 silence [88](#)
Destroy Clips command [196](#)
Destroy Selected Tracks (Remove and Close)
 command [261](#)
Device Order command [37](#)
Device Properties command [38](#), [39](#)
devices
 CD [262](#)
 choosing which devices to use [37](#)
 general uses of [37](#)
 setting properties for input [39](#)
 setting properties for output [38](#)
Dialogic ADPCM format [243](#)
DiamondWare Digitized format [243](#)

digital audio, fundamentals of [269](#)
Digital Theater Systems (DTS) [236](#)
DirectX
 plug-ins [36](#)
 using with surround sound [227](#)
disk space
 freeing up [53](#)
 monitoring [28](#)
Display Time Format command [63](#)
display window, navigating in [24](#)
Distortion effect [172](#)
dithering
 applying for playback [39](#)
 in response to lower bit depths [96](#)
 preferences [49](#)
 when saving 32-bit data to 16-bit files [51](#)
docking windows [23](#)
Dolby encoder [226](#)
Doppler Shifter effect [143](#)
downsampling
 audio during playback [39](#)
 quality level [50](#)
DTMF signals, generating [89](#)
DTS [236](#)
DTS encoder [226](#)
DVD, defined [276](#)
DVI/IMA ADPCM format [243](#)
DWD format [243](#)
Dynamic Delay effect [157](#)
Dynamic EQ effect [120](#)
Dynamic Peaks option [72](#)
dynamic range, displayed in Level Meters [71](#)
Dynamics Processing effect [133](#)

E

Early Reflections tab, Full Reverb effect [164](#)
EBU extensions [98](#)
echo
 adding to audio [156](#)
 defined [277](#)
Echo Chamber effect [159](#)
Echo effect [158](#)

Edit Favorites command [256](#)
Edit Original functionality, embedding in files [49](#), [222](#)
Edit Tempo command [216](#)
Edit View
 considerations for using [73](#)
 switching to [20](#)
 using [20](#)
 using effects in [138](#)
editing audio
 See also specific editing tasks
 general considerations [73](#)
 in Multitrack View [189](#)
 undoing and redoing [29](#)
effects
 See also specific effect names
 cancelling [51](#)
 general use of [33](#)
 graph controls [34](#)
 grouping in Effects tab [32](#)
 plug-ins [36](#)
 presets [33](#)
 real-time, about [205](#)
 using in Edit View [138](#)
Effects Rack dialog box [205](#)
Effects tab in Organizer window [32](#)
8-bit signed format [240](#)
empty audio clips, inserting [196](#)
Enable DirectX Effects command [36](#)
Enable Preroll And Postroll Preview command [36](#)
enhancing audio, about [101](#)
Envelope effect [132](#)
Envelope Follower effect [174](#)
envelopes
 defined [277](#)
 for clips [207](#)
equal loudness contour [251](#)
equalization (EQ), defined [277](#)
expander, defined [277](#)
exporting
 mixdowns to audio [239](#)

mixdowns to video [239](#)
Extract Audio From CD command [59](#)
Extract Audio From CD options [60](#)

F

fading audio files [130](#)
Fast Fourier Transform (FFT), defined [277](#)
fast-forwarding [70](#)
favorites
 about [256](#)
 viewing in Favorites tab [33](#)
feedback, about [146](#)
FFT Filter effect [122](#)
files
 See also specific format names
 assigning a bitmap image to [100](#)
 backward compatibility [49](#)
 choosing a format [240](#)
 closing [30](#)
 converting to a different format [237](#), [238](#)
 creating new [74](#)
 linking to Adobe Premiere Pro or After Effects [49](#), [222](#)
 monitoring size of [28](#)
 opening [56](#)
 properties, adding [98](#)
 renaming [253](#)
 saving audio to [237](#)
Files tab [30](#)
filtering audio, about [120](#)
Find Beats And Mark command [215](#)
Find Beats commands [78](#)
5.1 surround sound [226](#)
flange, defined [277](#)
Flanger effect [147](#)
Flush Virtual File command [54](#)
flushing
 defined [277](#)
 forcing complete flush [46](#)
Force New Row command [23](#)
frames, snapping to [79](#)

frequencies

 selecting [77](#)

 viewing [75](#)

Frequency Band Splitter effect [176](#)

frequency range, analyzing [102](#)

frequency, defined [277](#)

Full command [189](#)

Full Reverb effect [163](#)

Function tab, in Favorites dialog box [257](#)

FX Mixer window [206](#)

G

General Reverb tab, Full Reverb effect [164](#)

General tab, Advanced Session Properties dialog box [184](#)

Generate commands [89](#)

graph controls [34](#)

Graphic Equalizer effect [124](#)

Graphic Phase Shifter effect [150](#)

grid lines, showing or hiding [48](#)

Group Waveform Normalize command [250](#)

H

Hard Limiting effect [136](#)

Help [6](#)

Help tab, Favorites dialog box [258](#)

Hertz (Hz), defined [277](#)

hidden clips, revealing [196](#)

high frequency absorption time, about [168](#)

Hiss Reduction effect [114](#)

horizontal scroll bar [26](#)

Hybrid tool [182](#), [187](#)

I

IFF format [241](#)

importing

 audio [56](#)

 files [30](#)

 video files [223](#)

IMPS folder [170](#)

impulse

 and the Convolution effect [169](#)

 and the Full Reverb effect [163](#)

 defined [278](#)

Impulse (.imp) files [170](#)

Index cue

 about [83](#)

 defined [278](#)

input devices

 adjusting record levels [72](#)

 choosing [37](#)

 monitoring levels [71](#)

 setting properties [39](#)

input gain

 adjusting [72](#)

 monitoring [70](#)

Insert Audio From Video command [223](#)

Insert In CD Project command [260](#)

Insert Video command [223](#)

Insert/Delete Time command [181](#)

inserting

 empty audio clips [196](#)

 files into tracks [30](#)

interleaved 6-channel Windows Media 9 Pro (WMA) [234](#)

interpolate, defined [278](#)

inverting audio [89](#)

ISRC number [261](#)

K

keyboard shortcuts [21](#)

L

latency, correcting [40](#)

left channel

 editing [79](#)

 selecting data in [76](#)

Level Meters

 defined [278](#)

 using [70](#)

levels, adjusting [72](#)

LFE [229](#)

limiter, defined [278](#)
 limiting [251](#)
 Linear Energy Plot [48](#)
 linking sessions to Adobe Premiere Pro and After Effects [49](#)
 Load Meter [182](#)
 Logarithmic Energy Plot [48](#)
 Loop Duplicate command [192](#)
 Loop Info tab [217](#)
 Loop Properties command [218](#)
 loops
 about [214](#)
 defined [278](#)
 defining [215](#)
 extending or shortening [220](#)
 key of [216](#), [218](#), [219](#)
 previewing [69](#)
 properties in Edit View [216](#)
 properties in Multitrack View [218](#)
 stretching to fit [217](#)
 synchronizing [220](#)
 tempo of [216](#), [218](#), [219](#)
 lossless encoding [247](#)

M

macros [249](#)
 magnifying analysis graphs [106](#)
 magnifying the display [25](#)
 Manage Temporary Folder Reserve Space command [53](#)
 Marquee Selection tool [77](#)
 mastering
 audio [101](#)
 CDs [259](#)
 defined [278](#)
 memory, allocating buffers [45](#), [51](#)
 menus [21](#)
 Merge This Take (Destructive) command [67](#)
 Merge/Rejoin Split command [191](#)
 metadata [98](#)
 metronome [219](#)

Metronome tab, Advanced Session Properties dialog box [186](#)
 microphone placement, in Echo Chamber effect [160](#)
 Microsoft ACM [241](#)
 Microsoft DirectX [227](#)
 MIDI
 changing volume, pitch, and tempo [204](#)
 defined [278](#)
 fundamentals of [272](#)
 importing [203](#)
 mapping [203](#)
 playing [204](#)
 tracks, about [203](#)
 triggering [21](#)
 zooming [204](#)
 MIDI devices, choosing [37](#)
 MIDI Panic Button command [204](#)
 MIDI Timecode (MTC), defined [278](#)
 MIDI Trigger, defined [279](#)
 millisecond (ms), defined [279](#)
 miniplug, defined [279](#)
 Mix Down commands [213](#)
 Mix Down To CD Project command [260](#)
 Mix Gauge [182](#)
 defined [279](#)
 Mix Paste command [80](#), [95](#)
 mixdowns
 defined [279](#)
 exporting to audio [239](#)
 exporting to video [239](#)
 Mixers window [209](#)
 mixing
 automating with clip envelopes [207](#)
 default bit-resolution [52](#)
 defined [279](#)
 preferences [51](#)
 Mixing tab, Advanced Session Properties dialog box [185](#)
 m-Law Wave format [241](#)
 Monitor Record Level command [71](#)
 mono waveforms, converting to stereo [95](#)
 mono, defined [279](#)

- mouse
 - changing cursor [44](#)
 - zooming with mouse wheel [44](#)
- Move/Copy Clip tool [182](#), [187](#)
- moving windows [23](#)
- MP3 (ID3 Tag) file data [98](#)
- mp3Pro format [244](#)
- Multichannel Encoder
 - dialog box [226](#)
 - elements of [226](#)
 - Pan Envelope automation [230](#)
 - panning sound [228](#)
 - Preview Device menu [233](#)
 - Preview Format menu [233](#)
 - previewing [232](#)
- Multichannel Encoder command [227](#)
- multichannel WMA [235](#)
- Multitap Delay effect [161](#)
- Multitrack View
 - preferences [51](#)
 - switching to [20](#)
 - using [20](#)
 - work area of [180](#)
- multitrack-only effects [174](#)
- Music effect [173](#)
- muting
 - audio [87](#)
 - tracks [199](#)

N

- navigating in display window [24](#)
- new files, creating [74](#)
- New Session command [180](#)
- Next/Sun format [246](#)
- noise
 - generating [90](#)
- noise gate, defined [279](#)
- Noise Reduction effect [116](#)
- noise shaping
 - curve, for dithering [39](#)
 - defined [279](#)
 - when dithering [97](#)

- nondestructive editing, defined [279](#)
- Normalize effect [137](#)
- Normalize tab, Group Waveform Normalize dialog box [251](#)
- normalizing
 - defined [279](#)
 - files for CDs [259](#)
 - groups of files [250](#)
- Notch filter effect [125](#)
- Notes tab, Advanced Session Properties dialog box [187](#)
- Nyquist Frequency, defined [280](#)

O

- offline editing, defined [280](#)
- online Help [6](#)
- Open Append command [57](#)
- Open As command [57](#)
- Open Audio From Video command [223](#)
- Open command [56](#)
- Open Session command [58](#)
- opening files [56](#)
- order, defined [280](#)
- Organizer window [30](#)
- out-of-band peaks, working with [251](#)
- output devices
 - adjusting playback levels [72](#)
 - choosing [37](#)
 - setting properties [38](#)
- output gain
 - adjusting [72](#)
 - monitoring [70](#)
- overdrive [172](#)

P

- p.d.f. (probability distribution function) [39](#)
- pan envelopes [207](#), [230](#)
- Pan/Expand effect [140](#)
- Panner Point, in Multichannel Encoder [230](#)
- panning
 - in sessions [53](#)
 - in surround [228](#)

Parametric Equalizer effect [126](#)
pasting
 audio [80](#)
 resampling data [50](#)
pauses, adding between CD tracks [261](#)
pausing playing or recording [70](#)
PCM Raw Data format [248](#)
PCM, defined [280](#)
peak amplitude, displayed in Level Meters [70](#)
peak files
 defined [280](#)
 specifying options for [49](#)
performance, measuring with Mix Gauge and Load Meter [182](#)
phase
 analyzing [104](#)
 defined [280](#)
pink noise
 defined [280](#)
 generating [91](#)
Pitch Bender effect [151](#)
Pitch Correction Effect
 Automatic Correction mode [152](#)
 Manual Correction mode [153](#)
pitch, changing [151](#)
placekeeper windows [24](#)
play lists [86](#)
Play Looped button [215](#)
playback cursor [63](#)
playback devices. *See output devices*
playing audio
 about [67](#)
 automatically [56](#), [69](#)
 choosing a device for [37](#)
 downsampling [39](#)
 setting the start point [62](#)
plug-in
 defined [280](#)
 effects, using [36](#)
Power Indicator lines [229](#)

pre-emphasis, for CDs [262](#)
preferences [43](#)
premade loops [214](#)
preroll and postroll, for effect previews [35](#)
presets
 defined [280](#)
 using for effects [33](#)
previewing
 audio [56](#)
 audio in Organizer window [69](#)
 defined [281](#)
previewing effects
 buffer size for [45](#)
 in Edit View [35](#)
process signal, for Vocoder effect [177](#)
punch-in
 defined [281](#)
 deleting unused takes [52](#)
punching in audio [66](#)

Q

quantization, defined [281](#)
Quick Filter effect [127](#)
QuickVerb effect [166](#)

R

Radio Industry file data [98](#)
RAM, allocating buffers [45](#), [51](#)
ranges
 saving to new file [237](#)
 selecting in Multitrack View [182](#)
RAW format [240](#), [248](#)
RCA cable, defined [281](#)
real-time effects
 about [205](#)
 applying [205](#)
 locking tracks with [207](#)
 mixing [206](#)
 removing [205](#)
real-time, defined [281](#)

recording audio

32-bit options [39](#)

about [64](#)

choosing a device for [37](#)

correcting for drift [52](#)

setting the start point [62](#)

recording devices. *See input devices*

redoing edits [29](#)

referenced clip, defined [281](#)

reflective surfaces, simulating with Echo effect [158](#)

Refresh Effects List command [36](#)

Refresh Now command [182](#)

Remove All Tracks command [261](#)

Remove Clips command [196](#)

Remove Selected Tracks command [261](#)

removing noise [108](#)

renaming files [253](#)

repositioning windows [23](#)

resample, defined [281](#)

Rescale Volume Envelopes command [207](#)

reserve space [53](#)

resizing windows [23](#)

restoring audio, about [101](#)

restoring default workspace [44](#)

reverb

adding [163](#)

defined [281](#)

tails [167](#)

Reverb effect [167](#)

reversing audio [89](#)

Revert To Saved command [29](#)

rewinding [70](#)

ReWire tracks

mixing down [213](#)

ReWire tracks, working with [202](#)

right channel, editing [79](#)

right-click menus [21](#)

rip, defined [281](#)

ripping, from CD [59](#)

S

S/N ratio, defined [282](#)

SAM format [240](#)

sample rate

changing [94](#)

considerations for choosing [93](#)

defined [282](#)

previewing a different rate [94](#)

sample type

converting [56](#), [93](#)

displaying in status bar [28](#)

sample, defined [282](#)

sampler information, adding to files [99](#)

sampler, defined [282](#)

SampleVision format [246](#)

sampling rate, fundamentals of [270](#)

saving

audio files [237](#)

mixdowns to audio [239](#)

mixdowns to video [239](#)

sessions [238](#)

Scientific Filters effect [128](#)

.scp files [252](#)

script collection (.scp) files [252](#)

Script tab, in Favorites dialog box [258](#)

scripts

about [249](#)

working with [254](#)

Scripts command [254](#)

scrolling

about [26](#)

preferences in Edit View [44](#)

preferences in Multitrack View [53](#)

Select All Tracks command [260](#)

Select Entire Wave command [76](#)

selecting audio [76](#)

Selection/View Controls [27](#)

semitones, correcting with Pitch Correction effect [153](#)

separator bars [257](#)

sequencer, defined [282](#)

Session Properties command [219](#)

sessions

about [179](#)converting the sample rate [238](#)creating [180](#)default session [181](#)defined [282](#)deleting time in [181](#)inserting clips [58](#)inserting time in [181](#)linking to Adobe Premiere Pro or After Effects [222](#)opening [57](#)recording into [66](#)saving [238](#)using as SMPTE masters or slaves [183](#)Set Current Clipboard command [79](#)setting up Adobe Audition [37](#)Settings command [43](#)shortcuts [21](#)Show View Tabs command [21](#)

showing and hiding

toolbars [22](#)Transport Controls [64](#)view tabs [21](#)windows [24](#)shrinking audio [154](#), [195](#)

silence

deleting [88](#)generating [87](#)64-bit doubles (RAW) format [240](#)slip editing clips [189](#)smoothing edit boundaries [50](#)SMP format [246](#)

SMPTE

defined [282](#)monitoring synchronization [28](#)using sessions as masters or slaves [183](#)SMPTE Master Enable command [183](#)SMPTE Slave Enable command [183](#)SMPTE Start Offset command [183](#)Snap To Clips command [220](#)Snap to Frames command [223](#)Snap To Loop Endpoints command [220](#)Snap To Ruler (Coarse) command [220](#)Snap To Zero Crossings command [215](#)

snapping

about [78](#)clips [189](#)SND format [242](#), [246](#)soloing tracks [199](#)

sound card

defined [282](#)positioning information [45](#)sound wave, defined [282](#)sound, fundamentals of [267](#)

special effects

convolution [169](#)distortion [169](#)music [169](#)noise [90](#)tones [92](#)spectral display, adjusting [47](#)

Spectral View

adjusting display settings [48](#)selecting frequencies in [77](#)switching between Waveform View and [75](#)using [75](#)splines [34](#)Split command [191](#)SPTI (SCSI Pass Through Interface) [45](#)Static Peaks option [72](#)Statistics command [107](#)Status Bar window [27](#)

stereo

defined [282](#)imagery, changing [138](#)Stereo Chorus Mode [146](#)Stereo Expand [140](#)Stereo Field Rotate effect [141](#)

stereo waveforms

choosing which channel to edit [79](#)converting to mono [95](#)stop playing or recording [70](#)Stretch effect [154](#)stretching audio [154](#), [195](#)

stripe, defined [282](#)

Studio Reverb effect [168](#)

subwoofer [229](#)

Surround Panner [228](#), [229](#)

SVX format [241](#)

Sweeping Phaser effect [149](#)

symmetric dithering [50](#)

Synchronize Clips With Edit View command [62](#)

Synchronize Cursor Across Windows
command [62](#)

synchronizing

loops [220](#)

multitrack latency [40](#)

T

Take History command [67](#)

tempo

defined [282](#)

editing [64](#)

Tempo tab, Advanced Session Properties
dialog box [186](#)

temporary folders

managing size of [53](#)

setting location [46](#)

time display format [63](#)

Time Selection tool [182](#), [187](#)

time stretching clips [195](#)

Time window [63](#)

timecode, defined [282](#)

Timed Record Mode command [65](#)

timeline, snapping to ruler [79](#)

tones, generating [92](#)

Tool tab, Favorites dialog box [258](#)

toolbars [22](#)

tooltips [22](#)

touch tones, generating [89](#)

track controls, defined [283](#)

track cues

defined [283](#)

for CD [83](#)

inserting [259](#)

Track Equalizers window [200](#)

Track Mixer [209](#)

Track Properties window [197](#)

tracks

bit depth, setting [200](#)

changing volume [198](#)

channel, setting [200](#)

defined [282](#)

equalizing [200](#)

inserting files into [30](#)

locking [207](#)

muting [199](#)

naming [198](#)

panning [198](#)

soloing [199](#)

tracks for CDs

editing source audio [261](#)

inserting [259](#)

removing [261](#)

selecting [260](#)

setting properties [261](#)

Transport Controls

Fast Forward button [70](#)

Go to Beginning button [70](#)

Go to End button [70](#)

Pause button [70](#)

Play buttons [67](#)

Record button [65](#), [66](#)

Rewind button [70](#)

showing and hiding [64](#)

Stop button [70](#)

Trim command [82](#), [189](#)

Trim Digital Silence command [86](#)

tutorials [9](#)

TXT format [242](#)

U

undocking windows [23](#)

undoing edits

about [29](#)

preferences for [46](#)

unity gain, defined [283](#)

UPC/EAN, defined [263](#)

upsampling quality level [50](#)

V

Variable Bit Rate (VBR) [235](#)
VBR [235](#)
VBR encoding [244](#), [247](#)
vertical ruler
 changing scale of [75](#)
 scrolling in [26](#)
vibrato, and Chorus effect [146](#)
video
 clips [223](#)
 exporting mixdowns to [239](#)
 preparing mixdowns for export [225](#)
 previews, customizing [224](#)
 working with [222](#)
video files, importing audio and video from [223](#)
view tabs, showing and hiding [21](#)
views, switching between [20](#)
VOC format [243](#)
Vocoder effect [177](#)
volume envelopes [207](#)
volume, making consistent [250](#)
VOX format [243](#)
VST plug-ins [36](#)

W

WAV format [241](#), [243](#), [247](#)
Wave Cache [45](#)
wave file, defined [283](#)
Wave Properties command [98](#), [217](#)
waveform clip, defined [283](#)
waveform display
 defined [283](#)
 view modes for [74](#)
waveform statistics, viewing [107](#)

Waveform View

 adjusting display settings [48](#)
 switching between Spectral View and [75](#)
waveform, defined [283](#)
waveforms
 considerations for editing [73](#)
 fundamentals of [268](#)
 viewing [74](#)
wet, defined [283](#)
white noise
 defined [283](#)
 generating [91](#)
window handle [23](#)
Windows Media 9 Pro [226](#)
Windows Notepad [254](#), [256](#)
Windows PCM format [235](#), [247](#)
Windows Recording Mixer command [72](#)
windows, using [22](#)
WMA format [247](#)
workspace
 changing colors [47](#)
 changing display settings [48](#)
 described [19](#)
 restoring to default [44](#)
writing CDs [262](#)

Z

zero crossing, defined [283](#)
zero-crossing points
 adjusting selections to [77](#)
 automatically adjusting edits to [53](#)
 snapping to [79](#)
Zoom commands in Multichannel Encoder [232](#)
zooming [25](#)
zooming analysis graphs [106](#)