

ADOBE® AUDITION® 3

USER GUIDE



© 2007 Adobe Systems Incorporated. All rights reserved.

Adobe® Audition® 3.0 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 Premiere, After Effects, Creative Suite, and Photoshop are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries.

Dolby is a trademark of Dolby Laboratories. Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Apple, Mac OS, and Macintosh are trademarks of Apple Inc., registered in the United States and other countries. DEC is a trademark of Digital Equipment Corporation. Sun is a trademark or registered trademark of Sun Microsystems, Inc. in the United States and other countries. All other trademarks are the property of their respective owners.

MPEG Layer 3 audio compression technology is licensed by Fraunhofer IIS and Thomson.

mp3PRO audio coding technologies licensed from Coding Technologies, Fraunhofer IIS and Thomson Multimedia.

ReWire is a Propellerhead product.

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 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 CFR 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.

Adobe Systems Incorporated, 345 Park Avenue, San Jose, California 95110, USA.

Contents

Chapter 1: Getting started

Installation	1
Resources	2
New features	5

Chapter 2: Digital audio fundamentals

Understanding sound	8
Digitizing audio	10

Chapter 3: Workflow and workspace

Workflow	13
Viewing, zooming, and navigating audio	19
Customizing workspaces	23
Managing workspaces	29

Chapter 4: Setting up Adobe Audition

Connecting to inputs and outputs	31
Setting Adobe Audition preferences	33
Managing temporary files	41

Chapter 5: Importing, recording, and playing audio

Opening audio files in Edit View	43
Opening and adding to sessions in Multitrack View	46
Importing with the Files panel	48
Importing with Adobe Bridge	50
Importing audio from CDs	50
Getting ready for recording and playback	53
Recording audio	56
Playing audio	60
Monitoring recording and playback levels	62

Chapter 6: Editing audio files

Displaying audio in Edit View	65
Selecting audio	72
Copying, cutting, pasting, and deleting audio	76
Visually fading and changing amplitude	78
Working with markers	80
Creating and deleting silence	86
Inverting and reversing audio	87
Generating audio	88
Analyzing phase, frequency, and amplitude	91
Converting sample types	98
Recovery and undo	101

Chapter 7: Applying effects

Shared effects controls	102
Applying effects in Edit View	106
Applying effects in Multitrack View	107
Adding third-party plug-ins	110

Chapter 8: Effects reference

Amplitude and compression effects	112
Delay and echo effects	122
Filter and equalizer effects	128
Modulation effects	137
Restoration effects	141
Reverb effects	152
Special effects	158
Stereo imagery effects	162
Time and pitch manipulation effects	169
Multitrack effects	174

Chapter 9: Mixing multitrack sessions

Multitrack View overview	179
Basic track controls	182
Track routing and EQ controls	185
Arranging clips	191
Editing clips	195
Automating mixes	203
Synchronizing sessions with ReWire and SMPTE	208
Setting advanced session properties	210

Chapter 10: Composing with MIDI

MIDI sequencing basics	213
Importing, recording, and outputting MIDI	215
Displaying and editing MIDI	217

Chapter 11: Loops

Using loops in Multitrack View	221
Creating loops in Edit View	224

Chapter 12: Working with video

Working with video applications	227
Importing video and working with video clips	228
Previewing video	229

Chapter 13: Creating surround sound

Surround sound basics	232
Setting track pan and volume for surround mixes	233
Previewing and exporting surround mixes	236

Chapter 14: Saving and exporting	
Saving and exporting files	240
Audio file formats	244
Video file formats	252
Adding file information	253
Chapter 15: Automating tasks	
Favorites	256
Normalizing groups of files	258
Batch processing files	259
Scripts	261
Chapter 16: Building audio CDs	
Assembling CDs	264
Writing CDs	267
Chapter 17: Keyboard shortcuts	
Finding and customizing shortcuts	269
Default keyboard shortcuts	270
Chapter 18: Digital audio glossary	
Common audio terms	273
Index	280

Chapter 1: Getting started

Before you begin working with your software, take a few moments to read an overview of installation, online Help, and the many resources available to users. You have access to instructional videos, plug-ins, templates, user communities, and much more.

Installation

Requirements

❖ To review complete system requirements and recommendations for your Adobe® software, see the Read Me file on the installation disc.

Install the software

- 1 Close any other Adobe applications open on your computer.
- 2 Insert the installation disc into the disc drive, and follow the on-screen instructions.

Note: For more information, see the Read Me file on the installation disc.

Activate the software

If you have a single-user retail license for your Adobe software, you will be asked to activate your software; this is a simple, anonymous process that you must complete within 30 days of starting the software.

For more information on product activation, see the Read Me file on your installation disc, or visit the Adobe website at www.adobe.com/go/activation.

- 1 If the Activation dialog box isn't already open, choose Help > Activate.
- 2 Follow the on-screen instructions.

Note: If you want to install the software on a different computer, you must first deactivate it on your computer. Choose Help > Deactivate.

Register

Register your product to receive complimentary installation support, notifications of updates, and other services.

❖ To register, follow the on-screen instructions in the Registration dialog box, which appears after you install and activate the software.



If you postpone registration, you can register at any time by choosing Help > Registration.

Read Me

The installation disc contains the Read Me file for your software. (This file is also copied to the application folder during product installation.) Open the file to read important information about topics such as the following:

- System requirements

- Installation (including uninstalling the software)
- Activation and registration
- Font installation
- Troubleshooting
- Customer support
- Legal notices

Resources

Adobe Help resources

Documentation for your Adobe software is available in a variety of formats.

In-product and LiveDocs Help

In-product Help provides access to all documentation and instructional content available at the time the software ships. It is available through the Help menu in your Adobe software.

LiveDocs Help includes all the content from in-product Help, plus updates and links to additional instructional content available on the web. Find LiveDocs Help for your product in the Adobe Help Resource Center, at www.adobe.com/go/documentation.

Think of Help, both in the product and on the web, as a hub for accessing additional content and communities of users. The most complete and up-to-date version of Help is always on the web.

Adobe PDF documentation

The in-product Help is also available as a PDF that is optimized for printing. Other documents, such as installation guides and white papers, may also be provided as PDFs.

All PDF documentation is available through the Adobe Help Resource Center, at www.adobe.com/go/documentation. To see the PDF documentation included with your software, look in the Documents folder on the installation or content DVD.

Printed documentation

Printed user guides are available for purchase in the Adobe Store at www.adobe.com/go/store, where you can also find books published by Adobe publishing partners.

Extras

You have access to a wide variety of resources that help you make the most of your Adobe software. Some of these resources are installed on your computer during the setup process; additional helpful samples and documents are included on the installation or content disc. Unique extras are also offered online by the Adobe Exchange community.

Installed resources

During software installation, a number of plug-ins and presets are added to the application folder. Plug-ins provide effects for audio processing and VST instruments for MIDI sequencing. Presets include a collection of effect settings optimized for specific sonic results. If you don't want to create an effect from scratch, start with presets for inspiration.

Content on CD and DVD

On the installation disc, the Documentation folder contains a PDF version of Help, technical information, and other documents such as reference guides and specialized feature information.

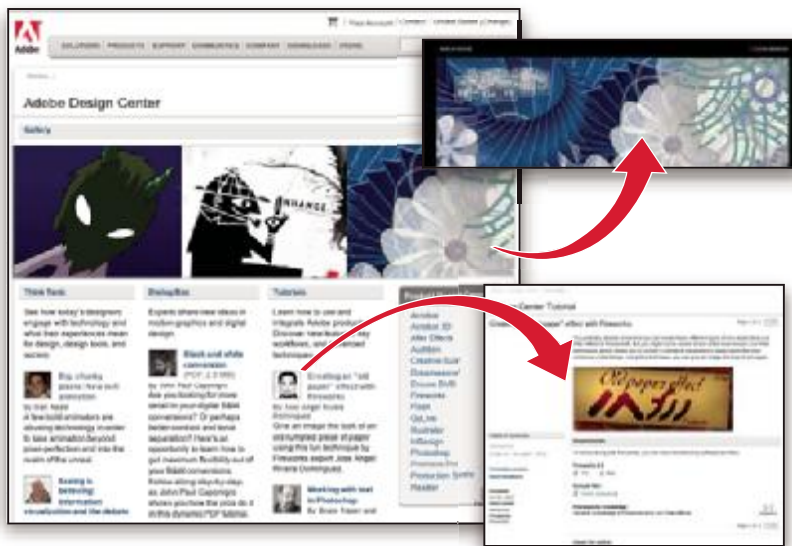
In addition to the installation disc, Adobe Audition® includes the Loopology DVD. This disc provides thousands of royalty-free loops, plus ready-made music beds and professional sample sessions.

Adobe Exchange

For more free content, visit www.adobe.com/go/exchange, an online community where users download and share thousands of free actions, extensions, plug-ins, and other content for use with Adobe products.

Adobe Design Center

Adobe Design Center offers articles, inspiration, and instruction from industry experts, top designers, and Adobe publishing partners. New content is added monthly.



You can find hundreds of tutorials for design products and learn tips and techniques through videos, HTML tutorials, and sample book chapters.

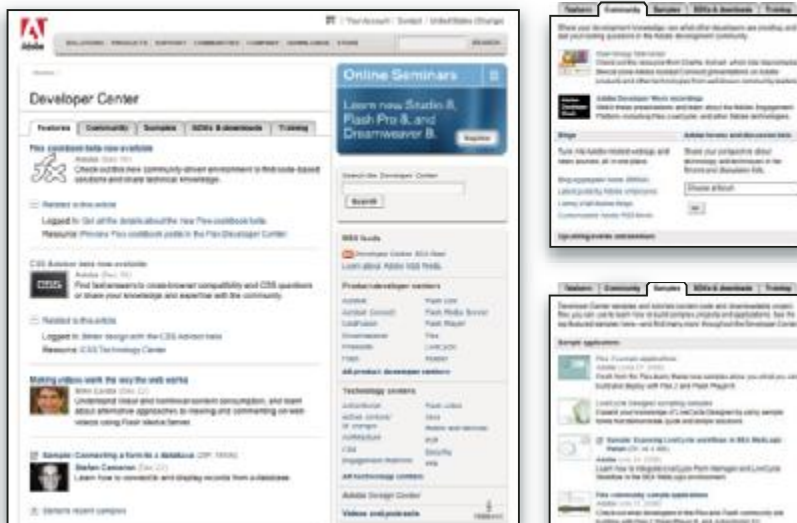
New ideas are the heart of Think Tank, Dialog Box, and Gallery:

- Think Tank articles consider how today's designers engage with technology and what their experiences mean for design, design tools, and society.
- In Dialog Box, experts share new ideas in motion graphics and digital design.
- The Gallery showcases how artists communicate design in motion.

Visit Adobe Design Center at www.adobe.com/designcenter.

Adobe Developer Center

Adobe Developer Center provides samples, tutorials, articles, and community resources for developers who build rich Internet applications, websites, mobile content, and other projects using Adobe products. The Developer Center also contains resources for developers who develop plug-ins for Adobe products.



In addition to sample code and tutorials, you'll find RSS feeds, online seminars, SDKs, scripting guides, and other technical resources.

Visit Adobe Developer Center at www.adobe.com/go/developer.

Customer support

Visit the Adobe Support website, at www.adobe.com/support, to find troubleshooting information for your product and to learn about free and paid technical support options. Click the Training link for access to Adobe Press books, a variety of training resources, Adobe software certification programs, and more.

Downloads

Visit www.adobe.com/go/downloads to find free updates, tryouts, and other useful software. In addition, the Adobe Store (at www.adobe.com/go/store) provides access to thousands of plug-ins from third-party developers, helping you to automate tasks, customize workflows, create specialized professional effects, and more.

Adobe Labs

Adobe Labs gives you the opportunity to experience and evaluate new and emerging technologies and products from Adobe.

At Adobe Labs, you have access to resources such as these:

- Prerelease software and technologies

- Code samples and best practices to accelerate your learning
- Early versions of product and technical documentation
- Forums, wiki-based content, and other collaborative resources to help you interact with like-minded developers

Adobe Labs fosters a collaborative software development process. In this environment, customers quickly become productive with new products and technologies. Adobe Labs is also a forum for early feedback, which the Adobe development teams use to create software that meets the needs and expectations of the community.

Visit Adobe Labs at www.adobe.com/go/labs.

User communities

User communities feature forums, blogs, and other avenues for users to share technologies, tools, and information. Users can ask questions and find out how others are getting the most out of their software. User-to-user forums are available in English, French, German, and Japanese; blogs are posted in a wide range of languages.

To participate in forums or blogs, visit www.adobe.com/communities.

New features

What's new

Record and mix

Adobe Audition 3.0 is a powerful tracking and mixing application. Mix faster with new automatic crossfades, clip fade handles, and automation-editing improvements. Take full advantage of the latest hardware with multicore processor support and an optimized mixing engine.

VST plug-in manager Quickly enable or disable specific VST plug-ins, optimizing performance. (See “Enable VST effects” on page 110.)

Auto crossfades and clip fade handles Simply overlap clips to crossfade them, and adjust fade curves with on-clip handles. (See “Fade or crossfade clips in a track” on page 197.)

Improved multitrack editing Efficiently edit sessions with these key enhancements:

- Collectively trim and fade grouped clips. (See “Grouping clips” on page 192.)
- Ripple-delete ranges of clips, instantly removing time gaps. (See “Trimming and extending clips” on page 195.)
- Adjust selected ranges of automation points. (See “Edit automation envelopes” on page 207.)
- Duplicate the contents of entire tracks, including clips, effects, and automation. (See “Duplicate a track” on page 184.)
- Simultaneously view all input and output levels to comprehensively monitor a mix. (See “Monitor levels” on page 63.)

XML session support Save sessions to XML format and other shared standards for multitrack applications. (See “Save multitrack sessions” on page 241.)

Mix down directly to Edit View Quickly output a session directly to Edit View, without first exporting a file. (See “Create a single audio clip from multiple clips” on page 201.)

Video previews for surround mixes Watch a preview in the Video panel while adjusting mixes in the Surround Encoder. (See “Previewing video” on page 229.)

Create and arrange

Adobe Audition 3.0 offers powerful and extensive looping capabilities, as well as support for VST instruments, making it easy to create and arrange great-sounding music. Improved processing, including the high-quality Radius time-stretching engine from iZotope and numerous new effects, gives you infinite creative options.

MIDI tracks and piano-roll editor Import, record, and edit MIDI, and output it through VST instruments or hardware synthesizers. (See Composing with MIDI.)

New effects Explore creative sonic possibilities with Convolution Reverb, Analog Delay, Guitar Suite, and other new effects. (See Effects reference.)

Radius time-stretching from iZotope Access industry-standard algorithms in the updated Stretch effect, as well as the File Info and Audio Clip Looping dialog boxes.

Bitmap audio images Export spectral graphs for detailed editing in an image-editor like Adobe Photoshop®. Or, import visually-oriented graphics as source material for experimental sound designs. (See “Spectral Bitmap Image (.bmp)” on page 250 and “Import a bitmap image as audio” on page 44.)

Improved CD ripping Automatically import track information from your favorite CD database. (See “Extract CD tracks with the Extract Audio From CD command” on page 51.)

Enhanced file sorting Sort files by track number, or by the date they were opened or created. (See “Change how files appear in the Files panel” on page 49.)

Customizable workspaces Tint panels and dialog boxes to suit your working style. Add favorite commands to the shortcut bar. (See “Change interface brightness or tint” on page 29 and “Display the shortcut bar” on page 27.)

Edit and master

Adobe Audition 3.0 includes a full set of editing, restoration, and mastering tools that give you unprecedented flexibility and control. Comprehensive waveform-editing tools combined with innovative spectral frequency brushes let you edit with power and precision. The new Mastering effect, phase correction tools, and Top/Tail view make Adobe Audition 3.0 the ideal audio editing and mastering environment.

Spot Healing Brush Quickly brush over artifacts to seamlessly remove them. (See “Select artifacts and repair them automatically” on page 74.)

Effects Paintbrush Create free-form selections, and layer brush strokes to determine the intensity of effects. (See “Select spectral ranges” on page 73.)

Marquee pan and phase selections Process discrete stereo information such as center-panned vocals in Spectral Pan Display or out-of-phase audio in Spectral Phase Display. (See “Select spectral ranges” on page 73.)

Play spectral selections Play back selected frequency, pan, and phase ranges to precisely restore and process audio. (See “Play audio linearly” on page 60.)

On-clip fade and gain controls Visually adjust selections or entire files. (See “Visually fading and changing amplitude” on page 78.)

Top/Tail View Fine-tune loop transitions by simultaneously viewing the beginning and end of files. (See “View the top and tail of an audio file” on page 71.)

Mastering effect Optimize audio for maximum impact with a series of professional processors. (See “Mastering effect” on page 161.)

Adaptive noise reduction Quickly correct a wide range of variable broadband noise. (See “Adaptive Noise Reduction effect” on page 141.)

Graphic Panner Visually adjust the stereo field to enhance spatial perception. (See “Graphic Panner effect” on page 165.)

Play lists Organize and play marker ranges for live performance and broadcast. (See “Creating play lists” on page 85.)

Efficient file opening and saving Specify default formats for Open and Save As dialog boxes, and quickly save groups of files to one format. (See “System preferences” on page 34 and “Save a group of audio files to one format” on page 240.)


Chapter 2: Digital audio fundamentals

Understanding sound

Sound fundamentals

Sound starts with vibrations in the air, like those produced by guitar strings, vocal cords, or speaker cones. These vibrations push nearby air molecules together, raising the air pressure slightly. The air molecules under pressure then push on the air molecules surrounding them, which push on the next set of molecules, and so on. As high-pressure areas move through the air, they leave low-pressure areas behind them. When these waves of pressure changes 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, it reflects these waves of air pressure. The zero line in the waveform is the pressure of air at rest. When the line swings up to a peak, it represents higher pressure; when the line swings down to a trough, it represents lower pressure.



*A sound wave represented as a visual waveform
A. Zero line B. Low-pressure area C. High-pressure area*

Waveform measurements

Several measurements describe waveforms:


Amplitude Reflects the change in pressure from the peak of the waveform to the trough. High-amplitude waveforms are loud; low-amplitude waveforms are quiet.

Cycle Describes a single, repeated sequence of pressure changes, from zero pressure, to high pressure, to low pressure, and back to zero.

Frequency Measured in hertz (Hz), describes the number of cycles per second. (For example, a 1000-Hz waveform has 1000 cycles per second.) The higher the frequency, the higher the musical pitch.

Phase Measured in 360 degrees, indicates the position of a waveform in a cycle. Zero degrees is the start point, followed by 90° at high pressure, 180° at the halfway point, 270° at low pressure, and 360° at the end point.


Wavelength Measured in units such as inches or centimeters, is the distance between two points with the same degree of phase. As frequency increases, wavelength decreases.



A single cycle at left; a complete, 20-Hz waveform at right
A. Wavelength B. Degree of phase C. Amplitude D. One second


How sound waves interact

When two or more sound waves meet, they add to and subtract from each other. If their peaks and troughs are perfectly *in phase*, they reinforce each other, resulting in a waveform that has higher amplitude than either individual waveform.




In-phase waves reinforce each other.


If the peaks and troughs of two waveforms are perfectly *out of phase*, they cancel each other out, resulting in no waveform at all.



Out-of-phase waves cancel each other out.

In most cases, however, waves are out of phase in varying amounts, resulting in a combined waveform that is more complex than individual waveforms. A complex waveform that represents music, voice, noise, and other sounds, for example, combines the waveforms from each sound together.

 *Because of its unique physical structure, a single instrument can create extremely complex waves. That's why a violin and a trumpet sound different even when playing the same note.*



Two simple waves combine to create a complex wave.

Digitizing audio

Comparing analog and digital audio

In analog and digital audio, sound is transmitted and stored very differently.

Analog audio: positive and negative voltage

A microphone converts the pressure waves of sound into voltage changes in a wire: high pressure becomes positive voltage, and low pressure becomes negative voltage. When these voltage changes travel down a microphone wire, they can be recorded onto tape as changes in magnetic strength or onto vinyl records as changes in groove size. A speaker works like a microphone in reverse, taking the voltage signals from an audio recording and vibrating to re-create the pressure wave.


Digital audio: zeroes and ones

Unlike analog storage media such as magnetic tape or 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 snapshots called *samples*. This process is typically known as *digitizing* or *sampling* the audio, but it is sometimes called *analog-to-digital conversion*.

When you record from a microphone into a computer, for example, analog-to-digital converters transform the analog signal into digital samples that computers can store and process.

Sample rate

Sample rate indicates the number of digital snapshots taken of an audio signal each second. This rate determines the frequency range of an audio file. The higher the sample rate, the closer the shape of the digital waveform is to that of the original analog waveform. Low sample 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 sample rate must be at least twice that frequency. (See “Nyquist frequency” on page 277.) For example, CDs have a sample rate of 44,100 samples per second, so they can reproduce frequencies up to 22,050 Hz, which is beyond the limit of human hearing, 20,000 Hz.


Here are the most common sample rates for digital audio:

Sample rate	Quality level	Frequency range
11,025 Hz	Poor AM radio (low-end multimedia)	0–5,512 Hz
22,050 Hz	Near FM radio (high-end multimedia)	0–11,025 Hz
32,000 Hz	Better than FM radio (standard broadcast rate)	0–16,000 Hz
44,100 Hz	CD	0–22,050 Hz
48,000 Hz	Standard DVD	0–24,000 Hz
96,000 Hz	High-end DVD	0–48,000 Hz

Bit depth

Just as sample rate determines frequency range, bit depth determines dynamic range. When a sound wave is sampled, each sample is assigned the amplitude value closest to the original wave's amplitude. Higher bit depth provides more possible amplitude values, producing greater dynamic range, a lower noise floor, and higher fidelity:


Bit depth	Quality level	Amplitude values	Dynamic range
8-bit	Telephony	256	48 dB
16-bit	CD	65,536	96 dB
24-bit	DVD	16,777,216	144 dB
32-bit	Best	4,294,967,296	192 dB



Higher bit depths provide greater dynamic range.

Audio file contents and size

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 requires 86 KB per second—about 5 MB per minute. That figure doubles to 10 MB per minute for a stereo CD, which has two channels.

 *In contrast to a digital audio file, a MIDI file might be as small as 10 KB per minute, so you can store up to 100 minutes of MIDI per megabyte. For more information, see “Understanding MIDI data and VST instruments” on page 213.*

How Adobe Audition digitizes audio

When you record audio in Adobe Audition, the sound card starts the recording process and specifies what sample rate and bit depth to use. Through Line In or Microphone In ports, the sound card receives analog audio and digitally samples it at the specified rate. Adobe Audition stores each sample in sequence until you stop recording.

When you play a file in Adobe Audition, the process happens in reverse. Adobe Audition sends a series of digital samples to the sound card. The card reconstructs the original waveform and sends it as an analog signal through Line Out ports to your speakers.

To summarize, the process of digitizing audio starts with a pressure wave in the air. A microphone converts this pressure wave into voltage changes. A sound card converts these voltage changes into digital samples. After analog sound becomes digital audio, Adobe Audition can record, edit, process, and mix it—the possibilities are limited only by your imagination.

Chapter 3: Workflow and workspace

Workflow

Workflow overview

Adobe Audition provides three distinct views, each optimized for unique audio workflows:


- In Edit View, you modify individual audio files, optimizing them for radio broadcast, the Internet, and audio CDs.
- In Multitrack View, you layer multiple audio files, mixing them together to create sophisticated musical compositions and video soundtracks.
- In CD View, you assemble audio files and convert them into tracks on audio CDs.

With the integrated environment of Adobe Audition, you can move seamlessly between these views, simultaneously editing and mixing files to create polished, professional audio. This integrated environment extends to Adobe video applications, where you can easily incorporate Adobe Audition into comprehensive video-editing workflows.

See also

“Comparing Edit View and Multitrack View” on page 20

Editing workflow



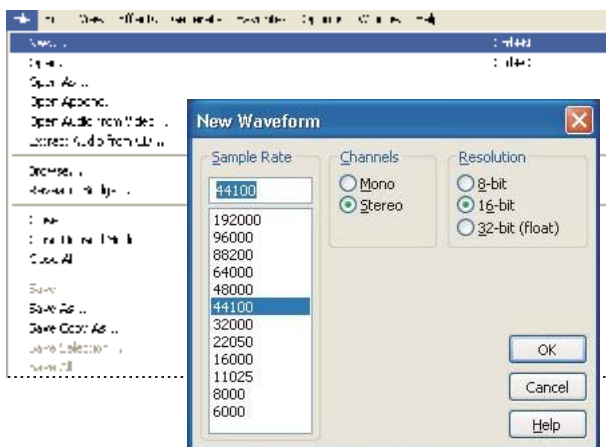
1. Open file 2. Edit audio 3. Apply effects 4. Save changes

In Edit View, you edit, restore, and enhance individual audio files, such as voiceovers, old vinyl recordings, and more. Any saved changes are permanent, making Edit View a great choice for mastering and finalizing files. (See “Editing audio files” on page 65.)

Open or create a file



Open an existing audio file that you want to modify. Alternatively, create a blank file that you’ll record or paste audio into. (See “Opening audio files in Edit View” on page 43.)

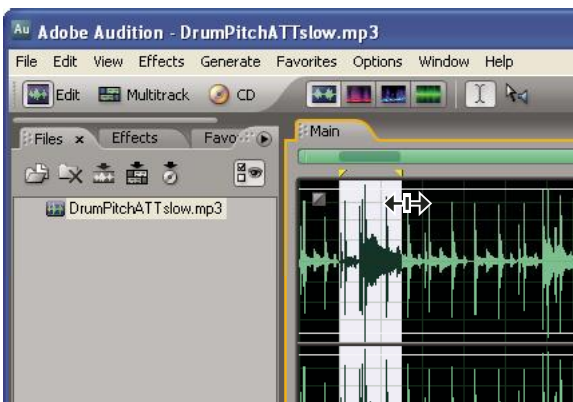


Creating a new file

Edit audio



In the Main panel, trim or extend a file by deleting or generating audio. To create a sonic collage, combine pasted audio from multiple files. Then, select noise or other audio you want to process with effects. (See “Selecting audio” on page 72.)

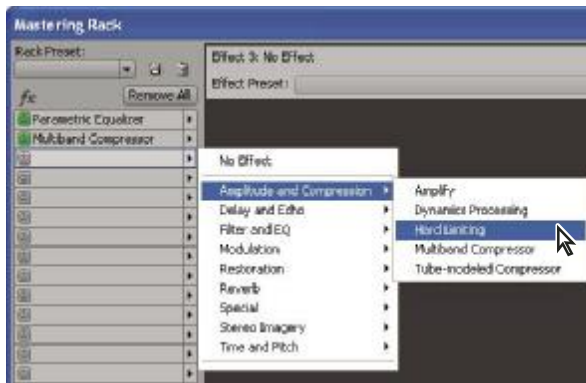


Editing audio in the Main panel

Apply effects



Apply process effects individually, or apply groups of VST effects in the Mastering Rack. In the rack, you can edit and reorder effects until you achieve the perfect results. (See “Applying effects in Edit View” on page 106.)

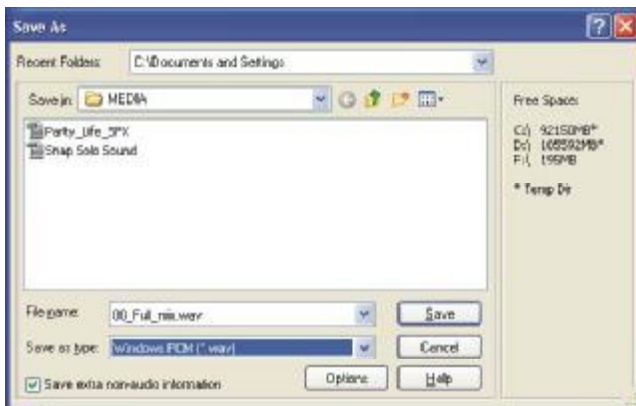


Applying groups of effects in the Mastering Rack

Save your changes




Save the polished, final file to disk, or automatically insert it in CD View or a multitrack session. (See “Saving and exporting files” on page 240 and Building audio CDs.)



Saving a file to disk

Multitrack workflow




1. Open session 2. Insert or record files 3. Arrange clips 4. Apply effects 5. Mix tracks 6. Export

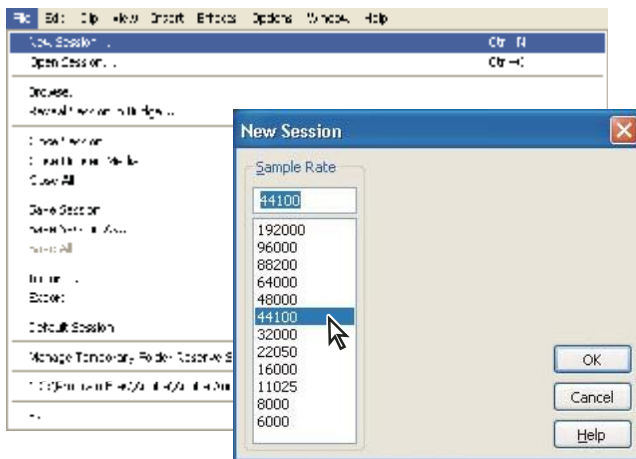
In Multitrack View, you layer multiple audio files to create stereo or surround-sound mixes. The edits and effects you apply aren’t permanent; if a mix doesn’t sound good next week, or even next year, simply change mix settings. (See “About multitrack sessions” on page 179.)

Open or create a session



From the File menu, choose Open Session or New Session. When you create a new session, you specify the sample rate for audio clips the session will contain. (See “Opening and adding to sessions in Multitrack View” on page 46.)

 To see how professional sessions are structured, open one of the sample sessions on the Adobe Audition Loopology DVD.

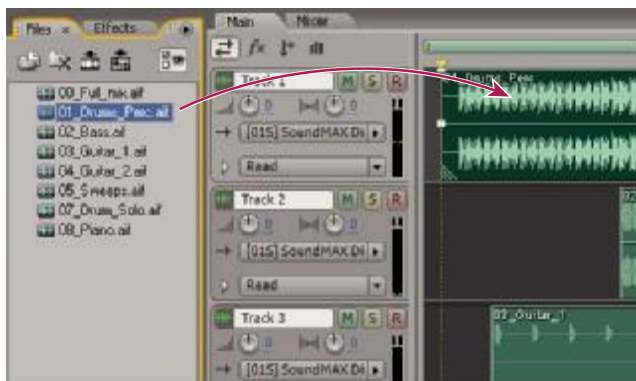


Creating a new session

Insert or record files



Insert audio, video, and MIDI files into tracks, or record new audio and MIDI while playing along with existing tracks. To build a particularly flexible session, insert audio loops—you can choose from over 5000 on the Adobe Audition Loopology DVD. (See “Insert an audio file into a session” on page 47 and “Record audio clips in Multitrack View” on page 57.)



Inserting from the Files panel

Arrange clips in the timeline



In the Main panel, arrange and edit clips in the timeline. In Multitrack View, edits are impermanent for maximum flexibility. But if you want to permanently edit a clip, simply double-click it to enter Edit View. (See “Arranging clips” on page 191 and “Editing clips” on page 195.)

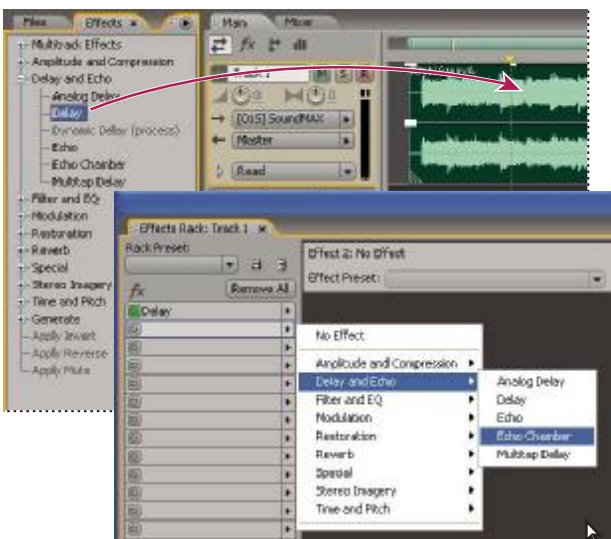


Arranging and editing clips in the Main panel

Apply effects



Apply effects in the Effects Rack, where you can edit, group, and reorder effects on each track. At any future time, you can update or remove effects to address the needs of different audio projects. (See “Applying effects in Multitrack View” on page 107.)



Applying effects in the Effects Rack

Route and mix tracks



In the Mixer or Main panel, output audio tracks directly to the Master track for standard mixes. As you build more complex mixes, combine related tracks in buses, and use sends to output individual tracks to multiple destinations. Then automate effect and mix settings over time, creating a dynamic, evolving mix that highlights different musical passages. (See “Track routing and EQ controls” on page 185 and “Automating track settings” on page 204.)

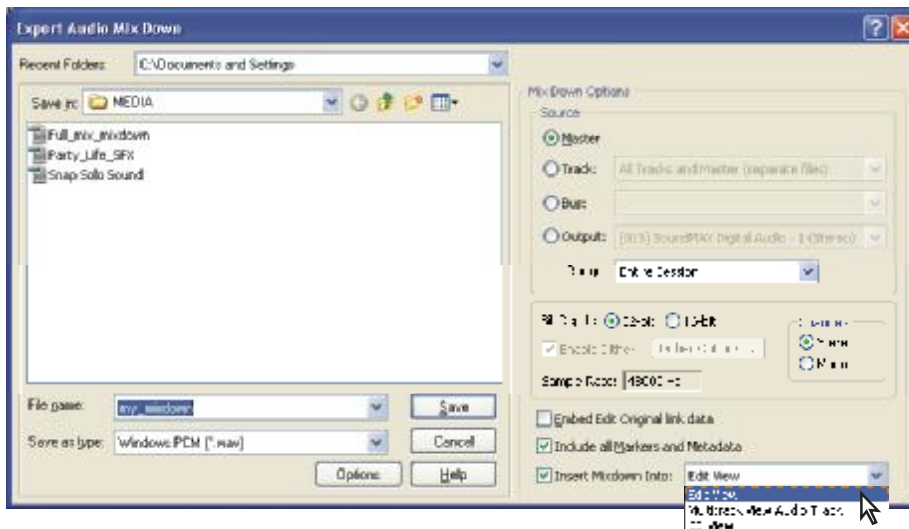


Routing and mixing tracks in the Mixer

Export



Export your finished mix to a file, which you can automatically insert in Edit View for mastering or CD View for archiving and distributing. (See “Export a session to an audio file” on page 242 and Building audio CDs.)



Exporting a mix to an audio file

Viewing, zooming, and navigating audio

Basic components of Edit, Multitrack, and CD View

The workspace differs in Edit View, Multitrack View, and CD View. However, all three views have similar basic components, such as view buttons, the Main panel, and the status bar.



Basic components of Edit, Multitrack, and CD View (Multitrack View shown)

A. View buttons B. Menu bar C. Toolbar D. Shortcut bar E. Main panel F. Various other panels G. Status bar

See also

“Comparing Edit View and Multitrack View” on page 20

“About CD View” on page 264

Comparing Edit View and Multitrack View

Adobe Audition provides different views for editing audio files and creating multitrack mixes. To edit individual files, use Edit View. To mix multiple files and integrate them with MIDI and video, 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 enter 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 about Edit View, see “Editing audio files” on page 65; for more information about Multitrack View, see Mixing multitrack sessions.

Switch views

❖ Do one of the following:

- From the View menu, choose Edit View, Multitrack View, or CD View.
- Click the Edit View , Multitrack View , or CD View button  in the toolbar.
- In Multitrack View, double-click an audio clip to open it in Edit View. Alternatively, double-click a file in the Files panel. Or, select an audio clip in either the Files panel or the Main panel, and then click the Edit File button  in the Files panel.

See also


“Comparing Edit View and Multitrack View” on page 20

“Keys for opening views” on page 270

Zooming audio

Zooming adjusts the view of the timeline display in the Main panel. The ideal zoom level depends on your current task. For example, you can zoom in to see details in an audio file or multitrack session, or you can zoom out to get an overview.

Adobe Audition provides several ways to zoom. You can either click buttons in the Zoom panel, or drag scroll bars and rulers.



Zooming methods

A. Click buttons in Zoom panel B. Drag scroll bars C. Right-click and drag rulers

See also

“Keys for playing and zooming audio” on page 270

Zoom with the Zoom panel

1 To display the Zoom panel, choose Window > Zoom Controls.

2 Do any of the following:

- Click the Zoom In Vertically button to increase the vertical resolution in Edit View or show fewer tracks in Multitrack View.
- Click the Zoom In Horizontally button to zoom in on the center of the visible waveform or session.
- Click the Zoom To Selection button to zoom in on the currently selected range.
- Click the Zoom In To Right Edge Of Selection button to zoom in on the right boundary of the currently selected range.
- Click the Zoom In To Left Edge Of Selection button to zoom in on the left boundary of the currently selected range.
- Click the Zoom Out Horizontally button to zoom out from the center of the visible waveform or session.
- Click the Zoom Out Full Both Axes button to display the entire audio file in Edit View or session in Multitrack View.
- Click the Zoom Out Vertically button to decrease the vertical resolution in Edit View or show more tracks in Multitrack View.

Zoom with a scroll bar or ruler


❖ In the Main panel, do any of the following

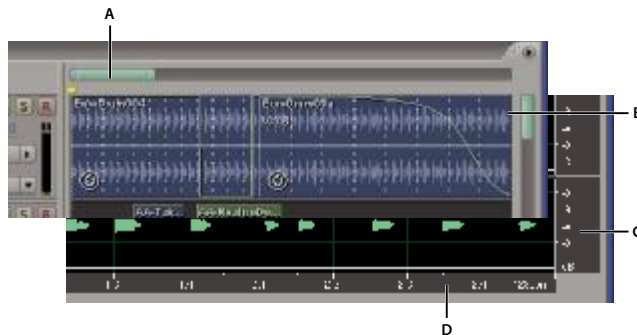
- Move the pointer to either edge of a horizontal or vertical scroll bar. When the pointer becomes a magnifying glass with arrows , drag left or right, or up or down.
- Right-click and drag in the horizontal ruler to zoom into a specific time range. The magnifying glass icon appears, creating a selection that shows you the range that will fill the Main panel.
- (Edit View only) Right-click and drag in the vertical ruler to zoom into a specific amplitude range. The magnifying glass icon appears, creating a selection of the range that will fill the Main panel.

To zoom with the mouse wheel, place the pointer over the appropriate scroll bar or ruler, and roll the wheel. (In Edit View, this zoom method also works when the pointer is over the waveform.) You can set the percentage of this zoom on the General tab of the Preferences dialog box. (See “General preferences” on page 33.)

Navigate with a scroll bar or ruler

At higher zoom levels, you may need to scroll to see different audio content in the Main panel.


 To change the placement of a horizontal scroll bar, right-click it, and choose *Above Display* or *Below Display*. This sets the position of the scroll bar for only the current view (Edit View or Multitrack View).



Scrolling devices

A. Horizontal scroll bar **B.** Vertical scroll bar (Multitrack View only) **C.** Vertical ruler (Edit View only) **D.** Horizontal ruler

- To scroll through time, drag the horizontal scroll bar, or drag left or right in the horizontal ruler.
- To scroll through audio amplitudes in Edit View, drag up or down in the vertical ruler.
- To scroll through tracks in Multitrack View, drag the vertical scroll bar.


 To scroll through tracks with the mouse wheel, place the pointer over the track display, and roll the wheel.

See also

“Navigate with the Selection/View panel” on page 22

Navigate with the Selection/View panel

The Selection/View panel shows the start, end, and length of the current selection and view in the Main panel. The panel displays this information in the current time format, such as Decimal or Bars And Beats. (See “To change the time display format” on page 55.)



Selection/View Controls panel

- 1 To display the Selection/View panel, choose Window > Selection/View Controls.
- 2 (Optional) Enter new values into the Begin, End, or Length boxes to change the selection or view.

See also

“Monitoring time during recording and playback” on page 54

“Dock, group, or float panels” on page 23

Customizing workspaces


About workspaces


Adobe video and audio applications provide a consistent, customizable workspace. Although each application has its own set of panels (such as Tools, Properties, Timeline, and so on), you move and group panels in the same way across products.

The main window of a program is the *application window*. Panels are organized in this window in an arrangement called a *workspace*. The default workspace contains groups of panels as well as panels that stand alone.

You customize a workspace by arranging panels in the layout that best suits your working style. You can create and save several custom workspaces for different tasks—for example, one for editing and one for previewing.

You can drag panels to new locations, move panels into or out of a group, place panels alongside each other, and undock a panel so that it floats in a new window above the application window. As you rearrange panels, the other panels resize automatically to fit the window.

 You can use floating windows to create a workspace more like those in previous versions of Adobe applications, or to place panels on multiple monitors.



Example workspace
A. Application window B. Grouped panels C. Individual panel



For a video about the Adobe workspace, see www.adobe.com/go/vid0249.

Dock, group, or float panels

You can dock panels together, move panels into or out of a group, and undock a panel so that it floats in a new window above the application window. As you drag a panel, *drop zones*—areas onto which you can move the panel—become highlighted. The drop zone you choose determines where the panel is inserted, and whether it docks or groups with other panels.

Docking zones



Docking zones exist along the edges of a panel, group, or window. Docking a panel places it adjacent to the existing group, resizing all groups to accommodate the new panel.



Dragging panel (A) onto docking zone (B) to dock it (C)

Grouping zones


Grouping zones exist in the middle of a panel or group, and along the tab area of panels. Grouping a panel stacks it with other panels.



Dragging panel (A) onto grouping zone (B) to group it with existing panels (C)


Dock or group panels

- 1 If the panel you want to dock or group is not visible, choose it from the Window menu.
- 2 Do one of the following:
 - To move an individual panel, drag the gripper area in the upper-left corner of a panel's tab onto the desired drop zone.



Drag panel gripper to move one panel

- To move an entire group, drag the group gripper at the upper-right corner onto the desired drop zone.



Drag group gripper to move entire group

The application docks or groups the panel, according to the type of drop zone.


Undock a panel in a floating window

When you undock a panel in a floating window, you can add panels to the window or otherwise modify it, as you do the application window. You can use floating windows to make use of a secondary monitor, or to create a workspace like those in earlier versions of Adobe applications.



- ❖ Select the panel you want to undock (if it's not visible, choose it from the Window menu), and then do one of the following:
 - Choose Undock Panel or Undock Frame from the panel menu. Undock Frame undocks the panel group.
 - Hold down Ctrl (Windows®) or Command (Mac OS®), and drag the panel or group from its current location. When you release the mouse button, the panel or group appears in a new floating window.
 - Drag the panel or group outside the application window. (If the application window is maximized, drag the panel to the Windows task bar.)

Resize panel groups


When you position the pointer over dividers between panel groups, resize icons appear. When you drag these icons, all groups that share the divider are resized. For example, suppose your workspace contains three panel groups stacked vertically. If you drag the divider between the bottom two groups, they are resized, but the topmost group doesn't change.

 To quickly maximize a panel beneath the pointer, press the tilde (~) key. (Do not press Shift.) Press the tilde key again to return the panel to its original size.

1 Do either of the following:

- To resize either horizontally or vertically, position the pointer between two panel groups. The pointer becomes a double-arrow .
- To resize in both directions at once, position the pointer at the intersection between three or more panel groups. The pointer becomes a four-way arrow .

2 Hold down the mouse button, and drag to resize the panel groups.




Dragging divider between panel groups to resize them horizontally
A. Original group with resize icon B. Resized groups

Open and close panels and windows

Even if a panel is open, it may be out of sight, beneath other panels. Choosing a panel from the Window menu opens it and brings it to the front.

When you close a panel group in the application window, the other groups resize to make use of the newly available space. When you close a floating window, the panels within it close, too.

- To open or close a panel, choose the panel from the Window menu.
- To close a panel or window, click its Close button .

Working with multiple monitors

To increase the available screen space, use multiple monitors. When you work with multiple monitors, the application window appears on the main monitor, and you place floating windows on the second monitor. Monitor configurations are stored in the workspace.

See also

“Dock, group, or float panels” on page 23

Display the toolbar

The toolbar provides quick access to tools, the Workspace menu, and buttons that toggle between Edit, Multitrack, and CD View. Some tools are unique to each view. Likewise, some Edit View tools are available only in spectral displays.

By default, the toolbar is docked immediately below the menu bar. However, you can undock the toolbar, converting it to the Tools panel, which you can manipulate like any other panel.

- To show or hide the toolbar, choose Window > Tools. A check mark by the Tools command indicates that it is shown.
- To undock the toolbar from its default location, drag the handle at the left edge to another location in the work area.
- To redock the Tools panel in its default location, drag the Tools panel tab to the drop zone that spans the entire width of the Adobe Audition window, just under the menu bar.



Available toolbar buttons differ in each view.

A. Edit View toolbar in Spectral Frequency Display B. Multitrack View toolbar

See also

“Dock, group, or float panels” on page 23

“Basic components of Edit, Multitrack, and CD View” on page 19

Display the shortcut bar

The shortcut bar displays buttons that provide quick access to commonly used functions. The shortcut bar appears in the upper part of the application window, below the menu bar and the default location of the toolbar.

The buttons contained in each shortcut group differ in Edit View, Multitrack View, and CD View. To identify a button, place the pointer over the button until a tool tip appears.

- To show or hide the shortcut bar, choose View > Shortcut Bar > Show.
- To show or hide shortcut groups, choose View > Shortcut Bar > Groups > [group name].
- To show or hide individual shortcuts, choose View > Shortcut Bar > [command type] > [command name].




You can display different shortcuts in Edit, Multitrack, and CD View.

See also

“Basic components of Edit, Multitrack, and CD View” on page 19

Display the status bar

The status bar runs across the bottom of the Adobe Audition work area. You can show or hide the status bar and select which types of information appear there.



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/MTC Slave Stability I. Display Mode

- To show or hide the status bar, choose View > Status Bar > Show. A check mark indicates that the status bar is visible.
- To change the type of information that the status bar displays, choose View > Status Bar or right-click the status bar, and select from the following options:

Data Under Cursor Shows information such as channel (for stereo files), amplitude (measured in decibels), and time (hours:minutes:seconds:hundredths of seconds) from the beginning of the audio file. This data changes dynamically when you move the pointer. For example, if you see R: -15.2 dB @ 0:00:242 when in Edit View, your pointer is over the right channel at 0.242 seconds, and the amplitude at that precise point is -15.2 dB.

In Multitrack View, you'll see even more data, such as envelope positions, dynamic effect settings, and the current position of clips as you drag them.

Sample Format Displays sample information about the currently opened waveform (Edit View) or session file (Multitrack View). For example, a 44,100 kHz 16-bit stereo file is displayed 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 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 time remaining 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 8-bit mono audio at 11,025 kHz, the time remaining might read 4399:15.527 free. Change the recording options to 16-bit stereo at 44,100 kHz, and the time remaining becomes 680:44.736 free.

In CD 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, helping you determine if these keys are fully pressed.

SMPTE/MTC Slave Stability Indicates the stability of incoming timecode compared to Adobe Audition's internal clock. For example, 95.0% SMPTE/MTC indicates a very strong SMPTE/MTC signal. Percentages above 80% should be stable enough to maintain sync. For more information, see "Synchronizing with SMTPE" on page 209.

Display Mode Indicates the current contents of the Main panel: Waveform, Spectral Frequency, Spectral Pan, or Spectral Phase Display in Edit View; Timeline in Multitrack View; CD List in CD View.

See also

"Basic components of Edit, Multitrack, and CD View" on page 19

Change interface brightness or tint

1 Choose Edit > Preferences, and click the Colors tab.

2 Do either of the following:

- To lighten or darken panels and dialog boxes, drag the UI Brightness slider. (To return to the original setting, click Reset UI.)
- To colorize panels and dialog boxes, click Tint, select or customize a color, and then click OK.

See also

“Colors preferences” on page 35

Managing workspaces

Choose a workspace

Each Adobe video and audio application includes several predefined workspaces that optimize the layout of panels for specific tasks. When you choose one of these workspaces, or any custom workspaces you’ve saved, the current workspace is redrawn accordingly.

❖ Open the project you want to work on, choose Window > Workspace, and select the desired workspace.

Predefined workspaces in Adobe Audition

Adobe Audition provides several predefined workspaces in the Window > Workspace submenu. These workspaces optimize the arrangement of panels for specific tasks.



You can customize any predefined workspace. (See “Save a custom workspace” on page 30.)

Edit View (Default) Arranges panels in the default positions for Edit View, and presents Waveform Display in the Main panel.

Multitrack View (Default) Arranges panels in the default positions for Multitrack View.

CD View (Default) Arranges panels in the default positions for CD View.

Frequency Space Editing Reflects the default arrangement for Edit View, but presents Spectral Frequency Display in the Main panel.

Mastering And Analysis Adds the Frequency and Phase Analysis panels to the default Edit View workspace.

Maximum Session (Dual Monitor) Arranges the Multitrack View work area for a two-monitor setup, displaying the Main panel and application window on one monitor and other panels on the second monitor, providing maximum view of the Main panel display.

Maximum Waveform Editing (Dual Monitor) Arranges the Edit View work area for a two-monitor setup, displaying the Main panel and application window on one monitor and other panels on the second monitor, providing maximum view of the Main panel display.

Session Mixing Adds the Mixer panel to the default Multitrack View workspace.

Video + Audio Session Adds the Video panel to the default Multitrack View workspace.

Save a custom workspace

As you customize a workspace, the application tracks your changes, storing the most recent layout. To store a specific layout more permanently, save a custom workspace. Saved custom workspaces appear in the Workspace menu, where you can return to and reset them.

❖ Arrange the frames and panels as desired, then choose Window > Workspace > New Workspace. Type a name for the workspace, and click OK.

Note: If a project saved with a custom workspace is opened on another system, the application looks for a workspace with a matching name. If it can't find a match (or the monitor configuration doesn't match), it uses the current local workspace.

Reset a workspace

Reset a workspace to return to its original, saved layout of panels.

❖ With the workspace you want to reset active, choose Window > Workspace > Reset *workspace name*.

Delete a workspace

1 Choose Window > Workspace > Delete Workspace.

2 Choose the workspace you want to delete, and then click OK.


Note: You cannot delete the currently active workspace.

Chapter 4: Setting up Adobe Audition

Connecting to inputs and outputs

About hardware inputs and outputs

You can use a wide range of hardware inputs and outputs with Adobe Audition. Sound card inputs let you bring in audio from sources such as microphones, tape decks, and digital effects units. Sound card outputs let you monitor audio through sources such as speakers and headphones. MIDI ports let you synchronize Adobe Audition with MIDI devices and applications.



A. Sound card inputs connect to sources such as microphones and tape decks. B. Sound card outputs connect to speakers and headphones. C. MIDI ports connect to MIDI devices and applications.

See also

“Synchronizing with ReWire” on page 208

Set audio inputs and outputs


When you set inputs and outputs for recording and playback, Adobe Audition can use two kinds of sound card drivers: Audio Stream In/Out (ASIO) and DirectSound. Some cards support both types of drivers.

ASIO drivers are preferable because they provide better performance and lower latency. You can also monitor audio as you record it and instantly hear volume, pan, and effects changes during playback. The main advantage of DirectSound is that you can access one card from multiple applications simultaneously.

- 1 Choose Edit > Audio Hardware Setup.
- 2 Click the Edit View, Multitrack View, or Surround Encoder tab.
- 3 For Audio Driver, choose a driver for the sound card you want to use. (Choose an ASIO driver if one is available; otherwise, choose the DirectSound driver, Audition Windows Sound.)
- 4 (Optional) Click Control Panel, set driver properties, and then click OK. For more information, see “Set driver properties for a sound card” on page 32.

5 Do any of the following:

- Under Edit View, choose stereo ports from the Default Input and Default Output menus.
- Under Multitrack View, choose stereo or mono ports from the Default Input and Default Output menus.

 *In a multitrack session, you can override the defaults for a particular track. (See “Assign audio inputs and outputs to tracks” on page 186.)*

- Under Surround Encoder, choose output ports for each surround channel in the Output Channel Mapping area. (For information about device requirements, see “Set the preview device for surround mixes” on page 237.)

Set driver properties for a sound card

To improve the performance of ASIO and DirectSound cards, optimize driver properties.

1 Choose Edit > Audio Hardware Setup, and click Control Panel.

2 Do one of the following, and then click OK.

- In an ASIO control panel provided by the sound card manufacturer, set driver options. (The options available will be different than those described below. For more information, consult the documentation for the sound card.)

Note: By default, Adobe Audition controls the ASIO sound card while playing or monitoring audio. If you want to access the card in another application, select Release ASIO Driver In Background.

- In the DirectSound Full Duplex Setup control panel, set the following options:

Device check boxes Use the check boxes under Direct Sound Output Ports and Direct Sound Input Ports to turn devices on or off. When a port is deselected, it is not available as a port option in the Audio Hardware Setup dialog box.

Buffer Size (Samples) Usually, the default settings work well for both playback (DirectSound Output Ports) and recording (DirectSound Input Ports), but some sound cards may require different settings. If you hear skips or dropouts in playback, you can adjust buffer size: Double-click the Buffer Size numerical entry for an input or output device, and type a new value.

Port Order If the selected device includes more than one port, click the Move Up or Move Down button to change the order of the ports for that device.

Sync Reference Specify whether you want the DirectSound Input or the DirectSound Output to set the master clock.

Full Duplex Select this to enable Adobe Audition to record an audio track while another one plays back, if the sound card is capable of doing so.

Start Input First This determines the order in which Adobe Audition starts the sound card playback (in) and record (out) ports in a multitrack environment. Select this option only if you have an older sound card that doesn't support full duplexing.

Connecting to MIDI devices and applications

Musical Instrument Digital Interface (MIDI) is a standard for communicating performance information from one piece of software or hardware to another. In Windows, you can transmit MIDI data internally between applications, or externally to and from devices such as MIDI keyboards. In Adobe Audition, you can utilize MIDI in the following ways:

- Using the MIDI sequencer, you can import, record, and edit MIDI data, and you can output it through VSTi virtual instruments or hardware synthesizers. (See Composing with MIDI.)

- Using a MIDI device connected to a MIDI input, you can trigger commands in Adobe Audition. For example, you can assign the Play command in Adobe Audition to the C4 note on your MIDI keyboard. (See “Enable MIDI-triggered shortcuts” on page 270.)
- Using MIDI outputs and inputs, you can send and receive SMPTE/MIDI timecode (MTC), synchronizing multi-track playback and recording with other hardware or software. (See “Synchronizing with SMPTE” on page 209.)

See also

“Synchronizing with ReWire” on page 208

Setting Adobe Audition preferences

Set Adobe Audition preferences

The Preferences dialog box lets you customize Adobe Audition’s display, editing behavior, use of memory and hard disk space, and other miscellaneous settings.

- 1 Choose Edit > Preferences.
- 2 Click a tab at the top of the dialog box to view a set of options.
- 3 When you finish setting options, click OK. To close the Preferences dialog box without changing any options, click Cancel.

When you click OK, most changes take effect immediately. If a change requires you to quit and restart Adobe Audition, you are prompted to do so. For example, you need to quit and restart Adobe Audition when you set up a different temporary folder.

See also

“General preferences” on page 33

“System preferences” on page 34

“Colors preferences” on page 35

“Display preferences” on page 36

“Data preferences” on page 37

“Multitrack preferences” on page 38

“SMPTE/MTC preferences” on page 40

“External Controllers preferences” on page 40

General preferences

In the Preferences dialog box, click the General tab to access the following options:

Force Spacebar To Always Trigger Play Plays a file when the spacebar is pressed, regardless of which dockable window has focus.

Auto-Play External Launch Lets you start Adobe Audition and play an audio file from the command line. (See “Play audio from the command line” on page 62.)

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

Upon A Manual Scroll/Zoom/Selection Change Determines auto-scrolling behavior when a manual scroll, zoom, or selection change occurs:

- **Abort Auto-Scrolling Until Next Play/Record** Stops auto-scrolling during the remainder of the current playback.
- **Resume Auto-Scrolling Only After Playback Cursor Enters View** Keeps the current portion of the track or waveform in view until the playback cursor reaches that point in the timeline, and then begins auto-scrolling to keep pace with the playback.
- **Resume Auto-Scrolling Immediately** Continues auto-scrolling whether or not the playback cursor is in view.

Edit View Right-Clicks Determines right-click behavior in the Main panel.

- **Popup Menu** Displays a menu. (To extend a selection, Shift-left-click.)
- **Extend Selection** Extends selected ranges instead of displaying the pop-up menu. (To see the pop-up menu, Ctrl-right-click.)

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

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.

Edit View Selections Determines the amount of waveform data that is automatically selected (if nothing is already highlighted) when you apply an effect.

- **View** Automatically selects only the area of the audio file that you can currently see on-screen.
- **Entire File** Automatically selects the entire waveform, even if you're only viewing a portion of it.



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

- **Select Audio After Paste** Automatically selects the inserted audio after a paste operation. Deselect this option to place the cursor at the end of the inserted audio instead.



Deselect this option to quickly append multiple segments of pasted audio.

Default Fade Curve Type Determines whether the default curve in Edit View is Linear/Logarithmic or Cosine. To enable the non-default curve, hold down Ctrl when dragging on-clip fade icons in the Main panel. (See “Visually fade in or out” on page 79.)

System preferences


In the Preferences dialog box, click the System tab to configure how Adobe Audition interacts with your system:

Cache Size Determines the amount of memory that Adobe Audition reserves for processing data. Recommended cache sizes range from 8 to 32 MB (32 MB is the default).

Undo Specifies behavior for the Edit > Undo command.

- **Enable Undo** Activates the undo function. Because undo requires extra disk space for temporary files and extra processing time, you may want to turn this feature off.

- **Levels (minimum)** Specifies the number of recent actions to keep when you purge undo. Lower this number to free up more memory but lose more actions.
- **Purge Undo** Deletes undo actions older than the minimum specified by the Levels setting, freeing up disk space but preventing you from reverting to previous edits.

Temporary Folders Lets you change the location of folders for temporary files, which Adobe Audition creates when you edit audio. (See “About temporary files” on page 41.) All temporary files begin with AUD and have a .tmp extension. Adobe Audition normally deletes temporary files when it exits. For the options below, click the Browse button  to navigate to a new folder location.

- **Primary Temp** Specifies the main folder for temporary files. Ideally, it should be on your fastest hard drive.
- **Secondary Temp** Specifies the secondary folder for temporary files. For best results, specify a different physical hard drive than the primary temp folder.

Note: If you have enough free space for temporary files on the drive that holds the primary temporary folder, you don't have to designate a secondary temporary folder.

CD Device Options Specifies the interface Adobe Audition uses to communicate with your CD drive: ASPI (Advanced SCSI Programming Interface) or SPTI (SCSI Pass Through Interface). Most CD drives support SPTI.

Auto-Save For Recovery Backs up files more frequently than needed for basic crash recovery. In Adobe Audition, basic crash recovery is on at all times. This option expands that feature, saving backups more frequently. Because this preference can affect performance, select it only if you experience frequent power failures or system crashes.

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

Force Complete Flush Before Saving Disables the quick save feature, which lets Adobe Audition quickly save files that contain only minor modifications. If you select this option, Adobe Audition saves an internal backup copy of entire files, considerably increasing the save time for large files. Select this option only if you have trouble saving back to the same filename or you have a problem with the Adobe Audition quick save feature.

Replace Invalid Filename Characters With This Character Replaces invalid filename characters in imported CD tracks. To delete invalid characters without replacing them, leave this box blank.

Default File Filter When Opening WAV Files Determines the variant of the WAV format used to open *.wav files. In most cases, Windows PCM is the best choice, as the vast majority of WAV files are stored in this format. Specify a different option only if you commonly open WAV files in a more unusual format.

Default Format For Open Dialogs Determines the default format for opening files. If you typically open files in one format, this option helps you narrow the displayed list of files. You can override the default format in the Open dialog box.

Default Format For Save As Dialogs Determines the default file format for saving files. If you typically save files in one format, this option makes the saving process more efficient. You can override the default format in the Save As dialog box.

Colors preferences

In the Preferences dialog box, click the Colors tab to change the Adobe Audition color scheme:

Waveform Lists all elements to which you can assign custom colors. To change the color of an element, select it from the list and click the Change Color button below the Example display.



To change clip color in Multitrack View, see “Set audio clip properties” on page 202.

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 pop-up menu.

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

Delete Deletes the currently selected color scheme preset.

Example Displays the currently selected color scheme preset or customized color scheme.

Change Color Opens the Color Picker, in which you can select a new color for the element selected under Waveform. (The current color is shown in the color swatch to the left of the Change Color button.)

Selection Lets you adjust the appearance of a selected range. Preview your changes in the Example area.

- **Invert** Completely reverses the color of selected ranges relative to unselected audio. Deselect this option if you want to set transparency for selected ranges instead of reversing the colors.
- **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. (This option is not available when Invert is selected.)

UI Brightness Drag the slider, click the arrows at either end of the brightness scale, or type in a percentage to change the overall brightness of the work area.

Tint Colors all panels and dialog boxes with a tint you specify.

Reset UI Returns UI Brightness and Tint to the default settings.

Display preferences

In the Preferences dialog box, click the Display tab to adjust the Spectral and Waveform Display modes:

Windowing Function Determines which method Adobe Audition uses to segment the spectral data before displaying 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 usually good choices.

Resolution Specifies the number of vertical bands used in drawing frequencies. The larger this number, the longer it takes for Adobe Audition to render the spectral display. Performance varies according to the speed of your computer.

Window Width Specifies the window width (FFT frame size) used to plot spectral data, where 100% is a full frame. The default setting is 75%. If you want to increase time resolution and determine exactly where frequencies start, decrease the window width (50% to 75% works best). This makes the display more accurate along the timeline (left and right), but less accurate along the frequency scale (up and down).

Decibel Range Adjusts the visible amplitude range for Spectral Frequency Display.

Initial Viewing Range Determines the default length of audio that the Top/Tail Views option displays.

Link Top/Tail Zoom Levels Automatically matches zoom levels in each view.

Show Marker And Range Lines Displays marker and range lines in the waveform display. Marker and range entries listed in the Marker List appear in the waveform as 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 Lines At option specifies the amplitude at which the boundary lines appear.

Peak Files Specifies options for peak (.pk) files, in which Adobe Audition stores 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 use 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** Saves peak (.pk) files in the same folder as all WAV files. You can safely delete peak files or deselect this option; however, without peak cache files, larger audio files reopen more slowly.
- **Rebuild Wave Display Now** Rescans the current file for sample amplitudes and redraws the waveform.

Show On-Clip Gain Control When There Is No Selection Range Always shows the on-clip control, so you can adjust the amplitude of an entire file without first making a selection.

Data preferences

In the Preferences dialog box, click the Data tab to control how Adobe Audition handles audio data:

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. All subsequent operations occur in the 32-bit realm.

Interpret 32-Bit PCM .wav Files As 16.8 Float Makes this version of Adobe Audition compatible with previous versions when handling 32-bit PCM WAV files.

Dither Transform Results Enables dithering when processing effects such as FFT Filter or Amplify. Adobe Audition does most processing using 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 fewer distortions and negative artifacts.

With this option enabled, you can approximate 24-bit performance with 16-bit data, because dithering increases dynamic range by about 10 dB. If this option is disabled, audio data is truncated to 16-bit during reconversion, and more subtle information is lost.

The drawback of dithering is that each operation adds a small amount of noise at the quietest volume levels. However, the trade-off between dithering (adding noise) and truncating the data (creating artifacts and limiting dynamic range) generally favors dithering, so it's best to select this option.

Use Symmetric Dithering Enables symmetric dithering. In most cases, it's best to select this option. If it is not selected, 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 Over Smooths cut and delete operations at the splicing point, preventing audible clicks at these locations. Applies only the boundaries over the number of milliseconds specified in the edit box.

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 end points, 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.

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.

- **Downsampling Quality Level** Enter a value from 30 to 1000. Higher values retain more high frequencies while still preventing the aliasing of higher frequencies to lower ones. A lower setting results in shorter processing time but also in the roll-off of certain high frequencies, 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 are fully satisfactory for most conversion needs. The default value is 80.
- **Prefilter** Prevents aliasing at downsampling. The prefilter removes all frequencies above the Nyquist limit, thus keeping them from generating alias frequencies at the low end of the spectrum. In general, select this option for best results.
- **Upsampling Quality Level** Enter a value from 30 to 1000. Higher values retain more high frequencies while still preventing the aliasing of higher frequencies to lower ones. A lower setting results in shorter processing time but also rolls off certain high frequencies, 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. Values between 100 and 400 address most conversion needs. The default value is 120.

 Use a higher Upsampling Quality Level 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 increased high-end filtering.

- **Postfilter** Prevents aliasing at upsampling. The postfilter 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 audio. The default value of 1 (bit) enables dithering, while a value of 0 disables dithering. For semidithering, choose a value of 0.5.

With dithering, you can approximate 24-bit performance with 16-bit data, because dynamic range increases by about 10 dB. This allows signals as quiet as -105 dB.

With this option enabled, you can approximate 24-bit performance with 16-bit data, because dithering increases dynamic range by about 10 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 preferences

In the Preferences dialog box, click the Multitrack tab to optimize recording, playback, and mix down:

Defaults Specifies recording and panning settings for multitrack sessions.

- **Recording Bit Depth** Select 16-bit or 32-bit.

- **Stereo Panning Mode** Select Left/Right Cut (logarithmic) to pan left by reducing the volume of the right channel, and pan right by reducing the left channel volume. The channel being panned to doesn't increase in volume as panning gets closer to 100%. Or select Equal-Power Sinusoidal to pan left and right channels with equal power, so that a hard pan left has 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%.

To override the default panning mode for a specific session, see "Mixing properties for sessions" on page 210.

Note: Because Equal-Power Sinusoidal panning can 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.

Audio Mix Down Contains options governing the resolution and dithering of multitrack sessions.

- **Bounce Bit Depth** Determines the bit resolution for the Edit > Bounce To Track command. Regardless of the session format (16-bit or 32-bit), you can bounce audio at either 16-bit or 32-bit quality with this option. The default is 16-bit.

- **Enable Dither** Select this option if the bit depth for mix down is lower than the bit depth of the session audio.

- **Dither Settings** Opens another dialog box with three options for dithering 16-bit mix downs: Dither Depth (in bits), Probability Distribution Function, and Noise Shaping. For more information, see "Change the bit depth of a file" on page 99.

Automation Optimizations Contains options that determine how automated mix parameters are recorded. (See "Techniques for automating mixes" on page 203.)

- **Automatch Time** Determines how fast parameters return to original values in the Touch or Latch modes. You can set values from 0.00 to 5.00 seconds; the default value is 1.00 second.

- **Linear Edit Point Thinning** Removes any edit points that represent static, unchanging parameter settings.

- **Minimum Time Interval Thinning** Creates edit points with intervals larger than the specified time value. Enter an interval between 1 and 2000 milliseconds in the Minimum Time field.

Auto Zero-Cross Edits Automatically adjusts the beginning and end points of all Cut, Copy, and Paste operations 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 end points 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.

Use BWF Time Reference As Insert Start Time Uses the timestamp embedded in a Broadcast Wave file (BWF) when inserting files. (See "Spot insert a Broadcast Wave file into a session" on page 47.)

Save Frozen Track Audio Files With Session Saves the temporary files associated with frozen tracks. When you reopen the session, Adobe Audition uses the temporary file instead of refreezing the tracks.

Play Audio While Scrubbing With The Start Time Indicator Plays back audio whenever you drag one of the yellow triangles for the Start Time Indicator, regardless of which tool is selected in the Tools panel. (See also "Play audio by scrubbing" on page 61.)

Crossfade Time Determines the crossfade duration for overlapping clips.

SMPTE/MTC preferences

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 "Synchronizing with SMPTE" on page 209.)

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.

In the Preferences dialog box, click the SMPTE/MTC tab to access options for timecode ports and synchronization.

SMPTE/MTC Output And Input The following options configure input and output ports for timecode:

- **MIDI Output Device (Send MTC To This Port)** Specifies the device to which Adobe Audition sends MIDI timecode.
- **MIDI Input Device (Receive MTC From This Port)** Specifies the device from which Adobe Audition receives MIDI timecode.
- **Use Internal Timestamps** Stamps incoming MIDI messages based on their arrival time in Adobe Audition, ignoring timestamps provided by Windows. (Some Windows systems provide uneven timestamps, causing imprecise synchronization.)

SMPTE/MTC Synchronization The following options apply only to incoming timecode.

- **Lead Time (To Prepare Audio Driver)** 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 (Freewheel - Audio Stops After No Input Detected For This Time)** Specifies the amount of time (in milliseconds) that Adobe Audition continues playing if it encounters a dropout in timecode.
- **Lag Time (Constant Time Difference Between Sound Card Driver's Reported Position And Actual)** 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 (Time Error Allowable Before Repositioning Play Cursor)** Specifies the number of frames that can fall out of sync with timecode before either repositioning the start-time indicator to match the code or performing a full resynchronization. A setting of up to 2.5 frames is recommended, because faulty timecode is usually corrected on the next frame sent. The default value is one-quarter (.25) of a frame.
- **Clock Drift Correction Time (Crossfade Length To Use When Adjusting Playback Position)** Specifies the number of samples to crossfade when making time corrections to chase audio to timecode. The default value is 250 samples.
- **Reposition Playback Cursor When Shuttleing** Readjusts the playback position if synchronization is off by the Slack value.
- **Full Re-Sync When Shuttleing** Performs a full re-synchronization if synchronization is off by the Slack value.

External Controllers preferences

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. With an external controller, you can automate settings for several tracks at once. (See "Automating track settings" on page 204.)

In the Preferences dialog box, click the External Controllers tab to access the following options:

External Control Device Specifies the controller you want to use.

Volume Increment Sets the minimum volume change for the data-entry knob on the controller.


Configure Accesses additional options provided by the controller software. Refer to your controller documentation for more information.

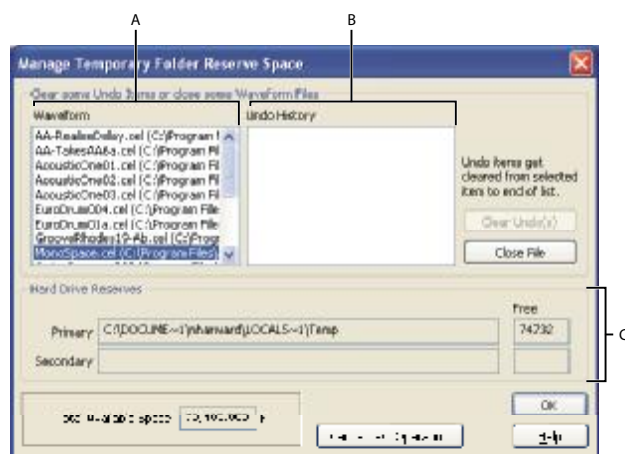
Managing temporary files

About temporary files

When you edit a file, Adobe Audition copies the audio data into a temporary file for internal use. This process allows for quicker editing, better handling of large files, and the ability to undo changes.

When specifying folders for temporary files, you should use hard disks with plenty of free space. (To specify these folders, see “System preferences” on page 34.) If free disk space reaches 45 MB, the Manage Temporary Folder Reserve Space dialog box appears. You can use this dialog box to delete temporary files and clear undo items.

 Use the status bar to monitor the amount of free disk space. (See “Display the status bar” on page 27.)



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


Manage temporary folder 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 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 (audxxxx.tmp) in your temporary folder that you should manually delete.*

See also

“System preferences” on page 34

Access an open file in another application

If an open Windows PCM (.wav) file hasn't yet been edited, and you need to access it in another application, flush a virtual version of the file. *Flushing* copies audio data from the file to a temporary folder. (This happens automatically if you alter the entire audio file.)

- ❖ In Edit View, choose File > Flush Virtual File.

Chapter 5: Importing, recording, and playing audio

Opening audio files in Edit View

Create a new, blank audio file

The File > New command creates blank audio files, which are perfect for recording new audio or combining pasted audio.

1 In Edit View, choose File > New.



To quickly create a new file from selected audio in an open file, choose Edit > Copy To New. (See “Copy or cut audio data” on page 76.)

2 In the New Waveform dialog box, set the following options, and click OK:

Sample Rate Determines the range of frequencies that can be reproduced. To reproduce a given frequency, the sample rate must be at least twice that frequency. (See “Sample rate” on page 10.)

Channels Determines if the waveform is mono or stereo.

- **Mono** Creates a waveform with just one channel of audio information. This option works well for a voice-only recording.
- **Stereo** Creates a two-channel waveform with separate right and left channels. This option is usually best for a music recording. However, because stereo waveforms contain twice as much data, they consume twice the storage space of mono waveforms.

Resolution Determines the range of amplitude levels that can be reproduced. The 32-bit level is best while you work in Adobe Audition. You can convert down for output later, if necessary. (See “Change the bit depth of a file” on page 99.)

Note: If you have an older sound card that can't play 32-bit files properly, you can use a lower bit depth for playback, without lowering the bit depth of a file. To check the capabilities of your sound card, choose Edit > Audio Hardware Setup, select the sound card under Audio Driver, and click Control Panel. For information about the available options, see your sound card documentation.

Open existing audio files

In Edit View, you can open audio from a variety of audio file formats, including MP3, WAV, and AIFF. You can also open the audio portion of video files in AVI, MPEG, MOV, or WMV format.

1 In Edit View, choose File > Open to open an audio file, or choose File > Open Audio From Video to open a video file.

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 (Optional) To preview the contents of a selected audio file, click Play to listen to the file once, select Loop to repeat the file until you click Stop, or select Auto Play to play files automatically when you select them.

4 Click Open.

See also

“About audio file formats” on page 244

“About video file formats” on page 252

Append one or more audio files to another

1 In Edit View, open the audio file that you want to add on to.

2 Choose File > Open Append.


3 In the Open Append dialog box, select one or more files. (To select multiple, adjacent files, click the first file and then Shift-click the last. To select multiple, nonadjacent files, Ctrl-click them.)

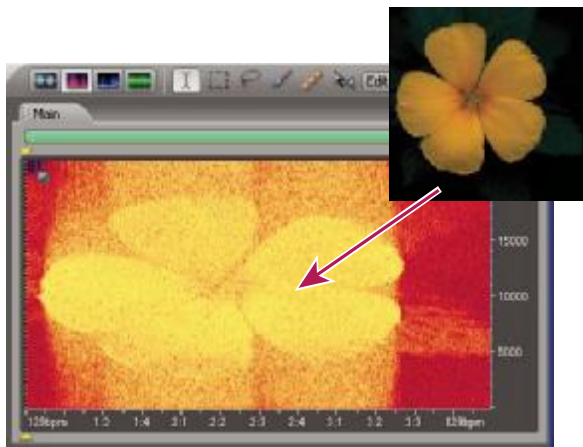
Note: If the selected audio files have a different sample rate, resolution, or channel type than the open audio file, Adobe Audition converts the selected files to match the open file. For the best results, append files with the same sample type as the original audio file.

4 Click Append.


Import a bitmap image as audio

In Edit View, you can import bitmap image (BMP) files created by any application. For the most predictable sonic results, import images of spectral graphs you saved from Adobe Audition and modified in an image-editing application. For sonic experimentation, import photographs, logos, and other visually-oriented files.

 When importing a saved spectral graph, make sure the related phase file is in the same folder. For more information about saving spectral graphs, see “Spectral Bitmap Image (.bmp)” on page 250.



Bitmap image imported into spectral frequency selection

1 (Optional) Open an audio file you want to incorporate an image into. Then, in Spectral Frequency Display , select a time or frequency range.

2 Choose File > Import Bitmap.

3 Select an image file in BMP format, and click Open.

4 Select one of the following import options:

Carve/Modulate Into Selection Applies the imported image as a mask, which filters existing spectral data.

Direct Paste Into Selection Completely replaces existing spectral data with the imported image.

Overlap Paste Into Selection Adds the imported image to existing spectral data, producing results similar to the Mix Paste command. (See “Mix audio data when pasting” on page 77.)

Full Import To New Waveform Creates a new audio file from the imported image. When you click OK, this option displays the New Waveform dialog box, where you specify sample rate, channels, and bit resolution. For more information about these options, see “Create a new, blank audio file” on page 43.

5 In the Advanced Options section, adjust the following settings if needed, and then click OK.

Note: Some options below are available only when specific import options are selected above.

White, Black Determine the decibel values that the brightest highlights and darkest shadows correspond to. Specify a small range to produce compressed audio or a large range to produce more dynamic audio.

Gamma Adjusts the midtone point, determining the relative amount of highlights and shadows. The standard gamma for Windows is 2.2. If the imported image was created in Mac OS, however, specify the standard gamma for that platform, 1.8. Or, specify more extreme gamma settings to produce experimental results.



If you set Black to a high amplitude value (such as -20 dB) and White to a low one (such as -100 dB), invert the gamma for better results. For example, instead of 2.0, use 0.5.

Frequency Spacing Available only with the Direct Paste and Overlap Paste options above, this setting determines how colors are distributed across the frequency spectrum. Lower values space colors more densely but require longer processing. Higher values provide quicker processing but reduce image detail. The default setting is 100 Hz for linear frequency display mode, or 100 Cents for logarithmic display mode. With cents, large values (50 or more) provide more musical results, particularly in Pure Tones mode.

Pure Tones Produces spectral lines at specific frequencies that vary in amplitude based on the image data. In logarithmic mode, harmonic results are easily achieved. In linear frequency display mode, results are very robotic and buzzy.

Random Noise Bands Adds a random frequency to each spectral line, smoothing out the harshness of pure tones. Solid color patches play as smoother noise, hiss, or rumble.

Track Spectral Features Achieves higher detail with fewer spectral lines. The results are less buzzy than with Pure Tones mode, with the ability to simulate noise cleanly as well. Individual lines in the source image are translated to individual frequency harmonics that follow the contour of the original image. Grass, whiskers, branches, and other features with lines are reproduced as spectral lines, providing a pen-and-ink or woodblock-print effect. Subtle patterns in the image are also preserved and audible. (By contrast, the two other modes above lose pattern detail.)

Log Amplitude or Linear Transfer Converts brightness values into audio amplitude either logarithmically (interpreting color data as decibels) or linearly (interpreting color data as amplitude percentages).

High Resolution Available only when Carve/Modulate or Full Import is selected above, this option increases image detail, requiring longer processing. High resolution is recommended if the bitmap image is an exported spectral graph that you modified in an image-editing application such as Adobe Photoshop.

Reset Returns all settings to the defaults, which are best for importing to a new file.

See also

“Saving and exporting files” on page 240

Convert audio to a different sample type when opening

- 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 Sample Rate, Channels, and Resolution options in the Open As dialog box, and click OK. For more information, see “Create a new, blank audio file” on page 43.

Opening and adding to sessions in Multitrack View

Create a new session

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 which files are a part of the session, where they are inserted, which envelopes and effects are applied, and so on.

- 1 Choose File > New Session.
- 2 Select a sample rate that your sound card supports. (To determine supported sample rates, consult the documentation for the card.)

Note: All files added to a session must share the sample rate. If you attempt to import files with different sample rates, Adobe Audition prompts you to resample them, which may lower audio quality.

- 3 (Optional) Select Use Default Session to use that session as a template. (See “Set or clear a default session” on page 47.)
- 4 Click OK.

See also

“About multitrack sessions” on page 179

“Sample rate” on page 10

“Save multitrack sessions” on page 241

Open an existing session

- 1 In Multitrack View, choose File > Open Session.
- 2 Locate and select the session (.ses) file, and click Open.

Note: Only one session can be open at a time.

See also

“Convert the sample rate of a session” on page 242

Set or clear a default session

After you set a default session, it opens whenever you start Adobe Audition. This default can also serve as a template for new sessions, letting you share settings, such as track inputs and volume levels, across multiple sessions.

Clearing a default session does not affect the session files themselves. It simply prevents Adobe Audition from automatically opening a specific session each time you start the application.

❖ In Multitrack View, do one of the following:

- To set a session as the default, open it, and choose File > Default Session > Set Current Session As Default. If the session contains clips, click Yes in the dialog box that appears.
- To specify no default session, choose File > Default Session > Clear Default Session.


Insert an audio file into a session

When you insert an audio file in Multitrack View, the file becomes an audio clip on the selected track. If you insert several files at once, or a single file that's longer than the space available on the selected track, Adobe Audition inserts new clips on the nearest empty tracks.

Note: For Broadcast Wave files, Adobe Audition can use the embedded timestamp to place files at the proper start times. (See “Spot insert a Broadcast Wave file into a session” on page 47.)

1 In Multitrack View, select a track, and then place the start-time indicator at the desired time position.

2 Do one of the following:

- To insert one or more audio files, choose Insert > Audio, select the files, and click Open. (To preview the contents of selected files, click Play or select Auto Play. Select Loop to repeat previews and Stop to end them.)
- To insert the audio portion of a video file in AVI, MOV, MPEG, or WMV format, choose Insert > Audio From Video, select a file, and click Open.
- To insert a currently open audio file, choose Insert > [file name].
- To insert a marker range from an open audio file, choose Insert > File/Marker List. In the window that appears, click the plus sign (+) to list a file's marker ranges, and select the range you want to add.
- To insert files from the Files panel, select them, and click the Insert Into Multitrack button , or press the Insert key.

See also

“About clips” on page 191

“Insert MIDI into a session and access the Sequencer” on page 214

“Insert a video file into a session” on page 228

“Use files in the Files panel” on page 49

Spot insert a Broadcast Wave file into a session


When you insert a Broadcast Wave (BWF) file into a multitrack session, Adobe Audition can use the embedded timestamp to insert the file at a specific time. This is commonly called *spot inserting*.

1 In Multitrack View, do any of the following to use embedded timestamps:

- Choose Insert > Insert Using BWF Timestamp.
- Right-click in the Files panel, and choose Use BWF Time Reference Offset.

- Click the Multitrack tab in the Preferences dialog box, and select Use BWF Time Reference As Insert Start Time.
- 2 In the Main panel, select a track.
 - 3 Choose Insert > Audio, and select a BWF file.

Adobe Audition inserts an audio clip at the designated start time.


 To view or edit the timestamp for a BWF clip, open the clip in Edit View, and then choose File > File Info. Click the Broadcast Wave tab. The timestamp value appears as the Time Reference.

See also

“Add audio file information” on page 253

Importing with the Files panel

Files panel overview

The Files panel displays a list of open waveforms, MIDI files, and video files for easy access. The Files panel also provides a variety of advanced options that let you show and hide markers, change the sort order of files, and play individual files. You can show or hide advanced options by clicking the Show Options button  in the upper right corner of the panel.




Files panel

See also

“Preview audio in the Files panel” on page 49

Import files into the Files panel

- 1 In the Files panel, click the Import File button .
- 2 In the Import dialog box, locate and select the file or files you want to import, and click Open. (You can select any audio, MIDI, or video files available on your computer or network.)

Note: For video files, the Files panel creates two clips: a video-only clip with the original file name (such as “clock.avi”) and an audio-only clip named “Audio for” followed by the original file name (such as “Audio for clock”).

See also

“Preview audio in the Files panel” on page 49

Use files in the Files panel




1 In the Files panel, select the files you want to work with:

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

2 At the top of the Files panel, do any of the following:

- To open a selected file in Edit View, click the Edit File button .


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

- To simultaneously insert each selected file into an individual track of the current session, click the Insert Into Multitrack button . (See “Insert an audio file into a session” on page 47.)
- To insert all selected files into CD View, click the Insert Into CD List button . (See “Insert CD tracks” on page 264.)
- To close all selected files, click the Close File button .







See also

“Preview audio in the Files panel” on page 49

Change how files appear in the Files panel

1 Click the Show Options button  at the top of the Files panel to display option buttons at the bottom of the panel.

2 Use the option buttons to do any of the following:





- To change the sort order of files, choose an option from the Sort By pop-up menu.
- To show or hide files according to type, select one or more of the following buttons: Show Audio Files , Show Loop Files , Show Video Files , and Show MIDI Files .
- To show or hide markers, click Show Markers . When Show Markers is selected, a plus-sign (+) appears next to any files that contain markers. Click the plus-sign to display the marker names. For more information, see “About markers” on page 80.
- To display the full path *[drive, folder(s), filename]* to each file, select the Full Path button .



Preview audio in the Files panel

The Files panel 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.

1 Click the Show Options button  at the top of the Files panel to display option buttons at the bottom of the panel.

2 Do any of the following:

- Select the file you want to preview, and then click the Play button . Click the Stop button  to stop the preview. Use the volume knob to adjust the volume of the preview.
- To automatically preview a file when you select it, click the Auto Play button .
- To enable continuous-loop preview, click the Loop Play button .
- (Multitrack View only) To preview loop-enabled files at the session tempo, select Follow Session Tempo.

 Loop-enabled files are identified with a loop icon .

See also

“Files panel overview” on page 48

Importing with Adobe Bridge

When you install Adobe Audition, you can optionally install Adobe® Bridge. Adobe Bridge is a cross-platform application that helps you locate, organize, and browse the assets you need to create audio, video, and web content.

From Adobe Bridge, you can:

- Preview, search, and sort audio and video files.
- Edit metadata for files.
- Place files into your audio projects.
- Perform automated tasks, such as batch commands.
- Start a real-time web conference to share your desktop and review documents.

For more details, either choose Help > Bridge Help in Adobe Bridge, or go to www.adobe.com/go/documentation.

Importing audio from CDs

About importing audio from CDs

If you want to import audio into Adobe Audition from a CD, you can digitally extract it or record it internally.

Digital extraction (also known as *ripping*) is the recommended method because it produces higher-quality audio than internal recording. Adobe Audition provides two methods for ripping tracks from CDs: the Open command and the Extract Audio From CD command. The Open command is the quickest method and is preferred for ripping entire tracks. The Extract Audio From CD command gives you more control, including the ability to rip partial tracks, access online CD databases, and customize the ripping process.

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 PCs 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 digitally from CD.

Extract CD tracks with 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 CD-ROM drive.
- 4 Select the tracks you want to rip, and click Open.

Extract CD tracks with the Extract Audio From CD command

- 1 Place an audio CD in the computer's CD-ROM drive.
- 2 In Edit View or CD View, choose File > Extract Audio From CD.
- 3 Under Device Interface Options, choose the drive that contains the audio CD. If necessary, set remaining interface options.
- 4 For Selection, do one of the following:
 - Select Absolute Block Selection to extract part of a track or a segment of audio that spans multiple tracks. In the Start box, enter the beginning frame. In the Length box, enter the total number of frames you wish to extract. (Each second of CD audio has 75 frames.) The resulting start and length times appear in Min:Sec:Frame format to the left of 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.



The option above is great for pulling hidden tracks from CDs, or for joining tracks that have been broken up by track indexes (such as live performances).

- 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. Select the desired tracks, and then set any of the following options:

Extract To Single Waveform Creates one audio file from all selected tracks.

Artist, Album, Genre, Year Adds track information you either manually enter or acquire automatically from a CD database.



To rename a track, double-click it.

Get Titles Acquires track information from a CD database you specify by clicking Configure.

Clear Removes all existing track information.

Configure Opens the Configure Titles dialog box, where you specify an online CD database and determine title format.

- 5 To listen to the selected tracks before extracting them, click Preview.
- 6 To save the settings for future use, save a preset. (See "Use effect presets" on page 104.)
- 7 After you finish setting options, click OK.

Configure track information from a CD database

- 1 In the Extract Audio From CD dialog box, click Configure.
- 2 In the HTTP Titles Database box, enter the domain name or IP address for your favorite CD database.
- 3 Enter your e-mail address. (Adobe Audition doesn't use this information, but many CD databases require it.)

4 In the Parse Song Title and Filename Template boxes, specify the information you want to include in the title and filename. Symbols correspond to information such as Artist (%a), Song Title (%s), Year (%y), and so on.

5 To capitalize first letters in words and trim spaces at the beginning and end of track titles, select the corresponding options.

6 Click OK.

Set Device Interface options

1 In the Extract Audio From CD dialog box, select 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 uses Input/Output control codes instead of SCSI commands.)

If you select ASPI/SPTI, set the following options:

Read Method Determines how 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** A SCSI 3-specific setting that works with most newer drives. If you have a newer CD-ROM drive, try this setting first.
- **SBC – Read(10)** A standard SCSI read setting that uses a 10-byte SRB (SCSI Request Block). All SCSI devices are required to support this setting.
- **SBC – Read(6)** 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.
- **NEC (D4)** 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 option deselected; select it only if the extraction process seems to work fine but the audio results contain only digital noise.

Swap Channels Places the left channel of CD audio in the right channel of the Main panel, and places the right channel of CD audio in the left channel of the panel.

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.

2 For Error Correction, CDDA Accurate is automatically selected if the CD-ROM drive has built-in ripping error correction. If your drive isn't CDDA Accurate, select either No Correction or Jitter Correction. (The latter compensates for data reading problems that older drives might have.)

Record from a CD internally

- 1 Preview the CD Audio input level to make sure that clipping won't occur. (See "Preview the CD Audio input level" on page 53.)
- 2 In Edit View, create a new file.
- 3 In the Transport panel, click the Record button .
- 4 Start the desired track in your CD player application.
- 5 When desired, stop recording in both Adobe Audition and the CD player application.

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 Edit View in Adobe Audition, and choose Options > Metering > Monitor Record Level.
- 3 Use the Levels panel 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 "Level meters overview" on page 62.)
- 4 If you need to adjust the CD Player input level, choose Options > Windows Recording Mixer to open the Windows Recording Control mixer. Adjust the CD Player input level as desired.
- 5 In Adobe Audition, deselect Options > Metering > Monitor Record Level.

Getting ready for recording and playback

Transport panel overview

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. Choose Window > Transport Controls to show or hide the Transport panel.



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



Transport panel

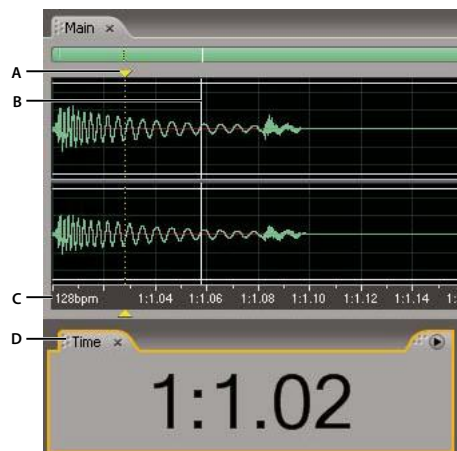
See also

"Playing audio" on page 60

Monitoring time during recording and playback

Adobe Audition provides several features to help you monitor time during recording and playback:

- The *start-time indicator* is a vertical, dotted yellow line in the Main panel. You set the start-time indicator to start playback or recording at a specific point in a waveform.
- The *playback cursor* appears as a vertical, solid white line in the Main panel. As you listen to audio files, the playback cursor moves, showing you the current time in the waveform.
- The Time panel shows the current time in numerical format. The default time format is Bars and Beats, but you can easily change it. The same format is used by the timeline along the bottom of the Main panel.



Features that help you monitor time

A. Start-time indicator B. Playback cursor C. Timeline D. Time panel

Set the start-time indicator

❖ Do one of the following:

- In the Main panel, click to move the start-time indicator to a new timeline position. (To perform this technique in Multitrack View, use any tool other than the Move/Copy tool.)
- In the Main panel, drag the handle (the yellow triangle) above or below the start-time indicator to a new timeline position.
- In the Selection/View panel, type a precise time in the Selection option under Begin, and press Enter.

After you set the start-time indicator, you can save it as a marker for later reference. See “About markers” on page 80.


Synchronize the start-time indicator across files or views


In Edit View, you can maintain the position of the start-time indicator when you switch between files—a useful technique when editing different versions of the same waveform. In Multitrack View, you can maintain the position of the start-time indicator when you switch to Edit View, a useful technique when applying edits and effects in both views.


- To synchronize the start-time indicator between waveforms in Edit View, choose Options > Synchronize Cursor Across Windows.
- To synchronize the start-time indicator between Multitrack View and Edit View, choose Options > Synchronize Clips With Edit View in Multitrack View.


Adjust the playback cursor


❖ Click one of the following buttons in the Transport panel:


Pause  Temporarily stops the playback cursor. Click the Pause button again to resume playback or recording.


Go To Beginning Or Previous Marker  Places the playback cursor at the beginning of the next marker. If there are no markers, the playback cursor moves to the beginning of the waveform or session.

Rewind  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.*

Fast Forward  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.*

Go To End Or Next Marker  Moves the playback cursor to the next marker. If there are no markers, the playback cursor moves to the end of the waveform or session.

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 “Calculate the tempo of a selected range” on page 225.

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.

Recording audio

Record audio in Edit View

You can record audio from a microphone or any device you can plug into the Line In port of a sound card. Before recording, you may need to adjust the input signal to optimize signal-to-noise levels. (See either “Adjust levels for a consumer sound card” on page 64 or the documentation for a professional card.)

- 1 Set audio inputs for Edit View. (See “Set audio inputs and outputs” on page 31.)
- 2 Do one of the following:
 - Create a new file.
 - Open an existing file to overwrite or append new audio, and place the start-time indicator where you want to start recording.
- 3 In the Transport panel, click the Record button  to begin recording. When you finish, click the Stop button .




To disable the Record button so that you don't start recording accidentally, right-click the Record button, and choose Disable Record Button. Repeat to reenable the button.

See also

- “Create a new, blank audio file” on page 43
- “Set the start-time indicator” on page 54
- “Record from a CD internally” on page 53
- “Monitoring recording and playback levels” on page 62

Use timed record mode in Edit View

- 1 Choose Options > Timed Record Mode. Alternatively, right-click the Record button in the Transport panel, and choose Timed Record Mode. A check mark indicates that Timed Record Mode is enabled.
- 2 In the Transport panel, click the Record button .
- 3 Select the maximum recording time:

No Time Limit Records until you click the Stop button (or until disk space runs out).

Recording Length Records for a limited period of time, which you type into the box, using the same display time format (such as Decimal or Bars And Beats) as the Main panel time display.

- 4 Select when to start recording:

Right Away Begins recording as soon as you click OK.

Time/Date Begins 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.

See also

- “Transport panel overview” on page 53

Direct-to-file recording in Multitrack View

In Multitrack View, Adobe Audition automatically saves each recorded clip directly to a WAV file. Direct-to-file recording lets you quickly record and save multiple clips, providing tremendous flexibility.

Prior to recording in a session, you must save it so that Adobe Audition can store recorded clips in the session folder. Inside that folder, you'll find each recorded clip in the *[session name]_recorded* folder. Clip file names begin with the track name, followed by the take number (for example, Track 1_003.wav).

After recording, you can edit takes to produce a polished final mix. For example, if you create multiple takes of a guitar solo, you can combine the best sections of each solo. Or, you can use one version of the solo for a video soundtrack, and another version for an audio CD.

See also


“Trimming and extending clips” on page 195

“Loop while recording in Multitrack View” on page 58

“Punch into a selected range in Multitrack View” on page 58

Record audio clips in Multitrack View

In Multitrack View, you can record audio on multiple tracks by overdubbing. When you overdub tracks, you listen to previously recorded tracks and play along with them to create sophisticated, layered compositions. Each recording becomes a new audio clip on a track.

1 In the Inputs/Outputs area  of the Main panel, choose a source from a track's Input menu.



Note: If you have a consumer sound card, you must also select the proper source in the Windows Recording Control mixer. (See “Adjust levels for a consumer sound card” on page 64.)

2 Click the Arm For Record button  for the track.

If you enable a track for recording in an unsaved session, the Save Session As dialog box appears. Specify a name and location, and click Save. (For more information, see “Direct-to-file recording in Multitrack View” on page 57.)

3 (Optional) To simultaneously record on multiple tracks, repeat steps 1-2 for each track.

4 In the Main panel, position the start-time indicator at the desired starting point, or select a range for the new clip.

5 In the Transport panel, click the Record button  to begin recording. When you finish, click the Stop button .

See also

“Set audio inputs and outputs” on page 31

“Monitoring recording and playback levels” on page 62

Hear inputs with effects applied to tracks


In Multitrack View, you can hear inputs with effects and sends that you've applied to a track, giving you an accurate preview of how recorded material will sound in the final mix. You can choose whether to hear an input all the time, or only when you reach a recordable portion of a track, such as a blank area or a punch-in range.

Note: ASIO sound cards are highly recommended for this feature. DirectSound cards increase latency (an audible delay) that disrupts timing for musicians. To reduce latency, see “Set driver properties for a sound card” on page 32.

1 Choose Options > Monitoring > Audition Mix, and then select one of the following:

Smart Input Monitors the input only when recording a track. When playing back, monitors only the clips on the track, not the input.

Always Input Monitors the input at all times during playback and recording. Any existing clips on the input track are not played.

2 To start monitoring an input, click an Arm For Record button  in the track controls.

3 To begin recording, click the Record button  in the Transport panel.

Hear inputs without effects applied to tracks

To avoid latency (an audible delay) when recording in Multitrack View, you may prefer to hear inputs directly from the sound card, without routing them through effects and sends in Adobe Audition.

1 Choose Options > Monitoring > External.


2 To set monitoring options, do one of the following:

- For a consumer sound card, use the Windows Volume Control mixer.
- For a professional sound card, see the card’s documentation.



Loop while recording in Multitrack View

If you loop while recording, Adobe Audition saves a new clip for each take. This feature is ideal for difficult musical passages. Simply loop while recording until the performer produces the perfect take. Or combine the best parts of each take to create a new clip. For more information, see “Display, remove, or play hidden clips” on page 195.

1 Set up one or more tracks for recording. (See “Record audio clips in Multitrack View” on page 57.)

2 In the Transport panel, right-click the Record button , 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 Loop While Recording button  to begin recording. When you finish, click the Stop button .

See also

“Transport panel overview” on page 53


“Direct-to-file recording in Multitrack View” on page 57

Punch into a selected range in Multitrack View


If you’re dissatisfied with a time range of a recorded clip, you can select that range and punch in a new recording, leaving the original clip intact. Though you can record into a specific range without punching in, punching in lets you hear audio immediately before and after a range; that audio provides vital context for creating natural musical transitions.


For particularly important or difficult sections, you can punch in multiple takes, and then select or edit takes to create the best performance. Adobe Audition saves two seconds of audio on either side of a punch-in range, so you can seamlessly edit and crossfade takes. (See “Display, remove, or play hidden clips” on page 195.)

Note: You can't punch into a loop-enabled clip. For information about disabling loops, see “Override original loop properties” on page 222.





A take created with the Punch In command

- 1 In the Main panel, drag the Time Selection tool  in the appropriate track to select a time range for the clip.
- 2 Make sure that the correct track input is selected. (See “Assign audio inputs and outputs to tracks” on page 186.)
- 3 Choose Clip > Punch In.
- 4 Position the start-time indicator a few seconds before the selected range.

The Arm For Record button  for the track is now activated in the Main panel.

- 5 In the Transport panel, do one of the following:

- To punch in a single take, click the Record button .
- To punch in multiple takes, right-click the Record button, choose a Loop While Recording option, and then click the Loop While Recording button .



See also

“Transport panel overview” on page 53

“Direct-to-file recording in Multitrack View” on page 57

Punch in during playback in Multitrack View


If you don't need to punch into a specific range, you can quickly punch into a general area during playback.

- 1 Enable one or more tracks for recording. (See “Record audio clips in Multitrack View” on page 57.)
- 2 In the Transport panel, click the Play button .
- 3 When you reach an area where you want to begin recording, click the Record button . When you finish recording, click the button again.

Playing audio



Play audio linearly

The Transport panel provides several options for playing audio linearly, from one time point to another. For example, you can play the visible portion of a file, the area from the start-time indicator to the end of the file, or the entire file.

 To start playback without using the Transport panel, press the space bar. Press the space bar again to stop playback.

1 In the Main panel, do either of the following:


- Set the start-time indicator.
- Select a range.

2 In the Transport panel, right-click the standard  or circled  Play button, and select one of the following options:




Play View Plays the visible portion of the file.

Play From Cursor To End Of View Plays from the start-time indicator to the end of the visible portion of the file. (This option is the default for the circled Play button.)

Play From Cursor to End Of File Plays from the start-time indicator to the end of the file. (This option is the default for the standard Play button.)

 To restrict playback to a selected time range, select one of the Play From Cursor options above.

Play Entire File Plays the complete file.

Play Spectral Selection Plays only frequencies you've selected with the Marquee , Lasso , or Effects Paintbrush  tool.

3 To start playback, click the Play button you customized.

Note: By default, the Main panel scrolls in sync with playback that extends beyond the visible section of a waveform. In the General tab of the Preferences dialog box, you can customize auto-scrolling options or disable this feature. (See “Set Adobe Audition preferences” on page 33.)

See also

“Transport panel overview” on page 53


“View the top and tail of an audio file” on page 71

“Play selections with preroll and postroll” on page 61

Play selections or files in a loop

If you need to fine-tune a selected range or experiment with different effects processing, loop playback.

1 If you want to loop a specific time range, select it in the Main panel.

2 In the Transport panel, right-click the Play Looped button , and choose either of the following:

Loop View (Or Selection) Plays the currently visible audio or selected range in a continuous loop.

Loop Entire (Or Selection) Plays the entire file or selected range in a loop.



3 To start playback, click the Play Looped button.

See also

“View the top and tail of an audio file” on page 71

Play selections with preroll and postroll

You can play selected ranges with *preroll* and *postroll* (short segments of audio immediately preceding and following the selection). Playing preroll and postroll lets you hear transitions so that you can fine-tune a selection.

1 In the Transport panel, right-click the standard Play button  or the circled Play button , and choose any of the following:

- Play Preroll And Postroll (Skip Selection)
- Play Preroll And Selection
- Play Postroll
- Play Preroll, Postroll, And Selection

2 To start playback in the selected preroll and postroll mode, click the Play button you customized.

3 (Optional) To deactivate preroll and postroll, right-click the standard or circled Play button, and choose any of the first four playback options.

See also

“Play audio linearly” on page 60

“Play selections or files in a loop” on page 60

Change preroll and postroll durations


By default, the duration of preroll and postroll is one second, but you can change these settings.


1 In Edit View, choose Options > Preroll And Postroll Options.

2 Under Edit View - Play, enter preroll and postroll durations for selected ranges in the Main panel.

3 Under Effects Preview, enter preroll and postroll durations for previews in effects dialog boxes.

Play audio by scrubbing

The term *scrubbing* comes from the process of manually rotating the reels of analog audio tape to play back audio. In Edit View and Multitrack View, you can use the Scrub tool  to achieve a similar result. Dragging this tool slowly plays audio at a slower rate; dragging it quickly increases playback speed up to the regular rate.

 *In Multitrack View, the Scrub tool scrubs individual tracks. To scrub all tracks, drag the start-time indicator.*

1 In the toolbar, select the Scrub tool .

2 In the Main panel, do one of the following:

- In Edit View, drag forwards or backwards over audio. (To scrub only one channel, drag at the top or bottom of the panel.)
- In Multitrack View, drag forwards or backwards over the track that you want to hear.

3 (Optional) Do any of the following:

- To create a selection while scrubbing, click to set the start-time indicator where you want the selection to begin, and then Shift-drag the Scrub tool.
- To simultaneously hear audio and move either edge of an existing selection, Shift-drag the Scrub tool.
- To play audio at the exact location of the pointer, Ctrl-drag the Scrub tool, or drag the start-time indicator.
- To play audio at a higher or lower speed than normal, Alt-drag the Scrub tool a short distance. Do not release the mouse button. The remaining audio plays at a speed comparable to the rate at which you dragged.

See also

“Select time ranges” on page 72

Play audio from the command line

By using the Run command in Windows, you can start Adobe Audition and begin playing a file in Edit View.

- 1** Choose Edit > Preferences, and select Auto Play On External Launch. Click OK.
- 2** In Windows, choose Start > Run.
- 3** Type the following text (including the quotation marks), and then click OK:

```
"[drive]:\Program Files\Adobe\Audition 2.0\Audition.exe" "[path to file]"
```

For example, to play the sound that Windows XP uses when you empty the Recycle Bin, type the following:

```
"c:\Program Files\Adobe\Adobe Audition 2.0\Audition.exe" "c:\Windows\Media\Windows XP Recycle.wav"
```


Note: Be sure to include one space between “. . . Audition.exe” and the next quotation mark.


Monitoring recording and playback levels

Level meters overview

To monitor the amplitude of incoming and outgoing signals during recording and playback, you use level meters. Edit View provides these meters only in the Levels panel. Multitrack View provides them in both the Levels panel, which shows the amplitude of an entire session, and track meters, which show the amplitude of individual tracks.

You can dock the Levels panel horizontally or vertically. When the panel is docked horizontally, the upper meter represents the left channel, and the lower meter represents the right channel.

 To show or hide the panel, choose Window > Level Meters.



Levels panel, docked horizontally
A. Left channel B. Right channel C. Peak indicators D. Clip indicators

The meters show signal levels 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 so that you can easily determine peak amplitude.

If amplitude is too low, sound quality is reduced; if amplitude is too high, clipping occurs and produces distortion. The red clip-indicator to the right of the meters lights up when levels exceed the maximum of 0 dB.



To clear clip indicators, either click them individually or choose Options > Metering > Reset All Indicators.

Monitor levels

Monitoring options differ in Multitrack View and Edit View. In both views, however, you can access these options from either the Options > Metering menu or the context menu that appears when you right-click meters.

Monitor levels in Multitrack View

Choose any of the following metering options:

Monitor Input Level Monitors audio inputs.

Meter Inputs Only Restricts monitoring to inputs. (By default, multitrack meters show both input and output levels.)

Show All Meters Displays input and output meters for all audio ports in the Levels panel. This option lets you quickly evaluate all levels in one place.



If your system has a large number of audio ports, you may need to resize the Levels panel to reveal meters for all of them.

Monitor levels in Edit View

Choose any of the following metering options:

Monitor Record Level Monitors audio inputs. (To quickly enable or disable this option, double-click the meters.)

Show Levels On Play And Record Enables the meters. (Disabling them improves performance on slower computers.)

Customize level meters

Either choose Options > Metering or right-click the meters. Then select any of the following options:

Range options Change the displayed decibel range.

Show Valleys Shows valley indicators at low-amplitude points.




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

Adjust For DC Compensates for DC offset. (For a definition, see “DC offset” on page 274.) During recording, any offset is indicated by clip indicators in the meters.

Note: Select *Adjust For DC* if your sound card records with a DC offset, shifting the center of the waveform above or below the zero amplitude line. This offset can dramatically shift the amplitude measured by meters, causing them to display levels inaccurately.

Dynamic or Static Peaks Change the mode of peak indicators. Dynamic Peaks resets the yellow peak level indicators to a new peak level after 1.5 seconds, letting you easily see recent peak amplitude. As the audio gets quieter, the peak indicators recede. Static Peaks retains peak indicators, letting you determine the maximum amplitude of the signal

since monitoring, playback, or recording began. However, you can manually reset peak indicators by clicking clip indicators.

 *To find out how loud audio will get before you record it, choose Static Peaks. Then monitor input levels and play the audio. After the audio ends, the peak indicators show the volume of the loudest part.*

Adjust levels for a consumer sound card

You may need to adjust levels if recordings are too quiet (causing unwanted noise), too loud (leading to clipped, distorted sound), or inaudible 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 peak somewhere between -2 dB and 0 dB when setting recording levels.


Adobe Audition doesn't directly control a sound card's recording and playback levels. For a professional sound card, you adjust these levels with the mixer application provided with the card. For consumer sound cards, however, you use the mixer provided by Windows.

1 Open the Windows Volume Control application, which resembles a small mixing board with vertical sliders.

You can usually access this application by choosing Start > All Programs > Accessories > Entertainment (or Multimedia). On many systems, you can also double-click the speaker icon in the system tray.

2 To adjust the sound card's playback (output) levels, drag Volume sliders as needed. (Below the sliders, ensure that Mute is deselected.)

3 To adjust the sound card's record (input) levels, choose Options > Properties, select Recording, and click OK. Select the input source you want to use, and adjust Volume sliders as needed.





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

Chapter 6: Editing audio files

Displaying audio in Edit View

Choose an audio display

In Edit View, you can choose any of four displays for audio data. Each shows you a different visual representation of sound waves. If you open a stereo file, the left channel appears at the top and the right channel appears at the bottom. If you open a mono file, its single channel fills the total height of the Main panel.

❖ In the toolbar, select Waveform Display , Spectral Frequency Display , Spectral Pan Display , or Spectral Phase Display .

See also


“Sound fundamentals” on page 8

“Comparing Edit View and Multitrack View” on page 20

About Waveform Display

Waveform Display shows a waveform as a series of positive and negative peaks. The x-axis (horizontal ruler) measures time, and the y-axis (vertical ruler) measures amplitude—the loudness of the audio signal. Quiet audio has both lower peaks and lower valleys than loud audio. You can customize Waveform Display by changing the vertical scale, colors, and line visibility.

With its clear indication of amplitude changes, Waveform Display is perfect for identifying percussive changes in vocals, drums, and more. To find a particular spoken word, for example, simply look for the peak at the first syllable and the valley after the last.



Stereo file in Waveform Display

See also

“Change the vertical scale” on page 71

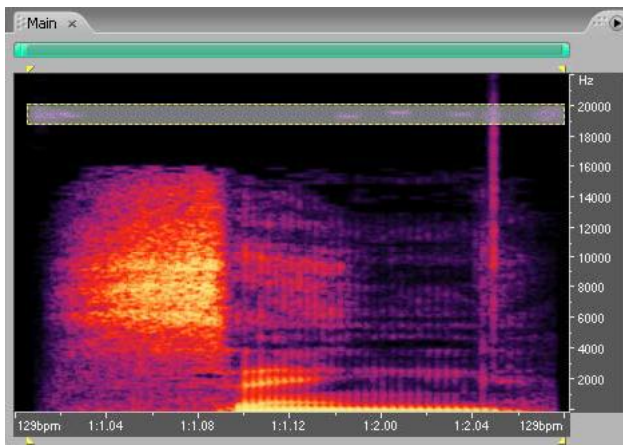
“Colors preferences” on page 35

“Display preferences” on page 36

About Spectral Frequency Display

Spectral Frequency Display shows a waveform by its frequency components, where the x-axis (horizontal ruler) measures time and the y-axis (vertical ruler) measures frequency. This view lets you analyze audio data to see which frequencies are most prevalent. Brighter colors represent greater amplitude components. Default colors range from dark blue (low-amplitude frequencies) to bright yellow (high-amplitude frequencies). You can modify display settings for Spectral Frequency Display using the Spectral Controls panel.

Spectral Frequency Display is perfect for removing unwanted sounds, such as coughs and other artifacts. This is known as *frequency-space editing*.



Spectral Frequency Display, with high frequencies selected

See also

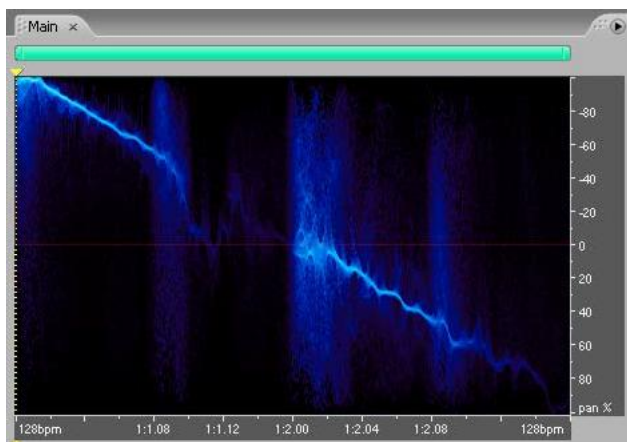
“Select spectral ranges” on page 73

“Adjusting spectral displays” on page 68

“Techniques for restoring audio” on page 141

About Spectral Pan Display

Spectral Pan Display shows the pan (left-right stereo) position of every frequency in an audio file, so you can visualize the sound locations in the stereo plane. The x-axis (horizontal ruler) measures time and the y-axis (vertical ruler) measures pan position, with the top of the display (-100%) representing full left and the bottom of the display (100%) representing full right. Brighter colors represent louder audio signals. You can modify display settings for Spectral Pan Display using the Spectral Controls panel.



Spectral Pan Display, with audio moving from left to right

Spectral Pan Display works well in conjunction with the Center Channel Extractor effect. That effect lets you isolate frequencies that are common to both the left and right channels (in other words, sounds that are panned center). In some cases, you may find that the audio you want to isolate is slightly off-center. You can use Spectral Pan Display to determine the actual pan position of the audio you want to isolate and enter that value in the Center Channel Extractor to improve the quality of the extraction. You can also use the Center Channel Extractor to boost the signal of audio at a particular pan position. For more information, see “Center Channel Extractor effect” on page 163.

Another effect that works well with Spectral Pan Display is the Pan/Expand effect. The Pan/Expand effect lets you shift the center channel of a stereo waveform and expand or narrow the stereo separation of the left and right channels. For more information, see “Pan/Expand effect (Edit View only)” on page 166.

See also

“Adjusting spectral displays” on page 68


“Noise Reduction effect (Edit View only)” on page 149

About Spectral Phase Display

Spectral Phase Display shows the phase difference, in degrees, between the left and right channels. For example, if any frequency is 180° out of phase, Spectral Phase Display shows bright patches near the +/-180° marks. You can modify display settings for Spectral Phase Display using the Spectral Controls panel.



Audio that is more than 90° out of phase creates problems when summing to mono and may sound odd in stereo. To help you determine how much audio is out of phase, Adobe Audition displays lines at the 90° marks by default.



Spectral Phase Display, with audio moving from out-of-phase (at top and bottom) to in-phase (at center)

For maximum precision, use Spectral Phase Display in conjunction with the Phase Analysis panel. (See “Analyze phase” on page 91.) For example, if you see a lot of inverse phase information in the Phase Analysis panel, you can use Spectral Phase Display to find out exactly where in time the out-of-phase audio is occurring.

To correct a phase problem, you can use the Effects > Invert command to inverse phase by 180° or the Center Channel Extractor to extract audio with a certain phase relationship. For more information, see “Invert a waveform” on page 87 and “Center Channel Extractor effect” on page 163.

See also

“Adjusting spectral displays” on page 68

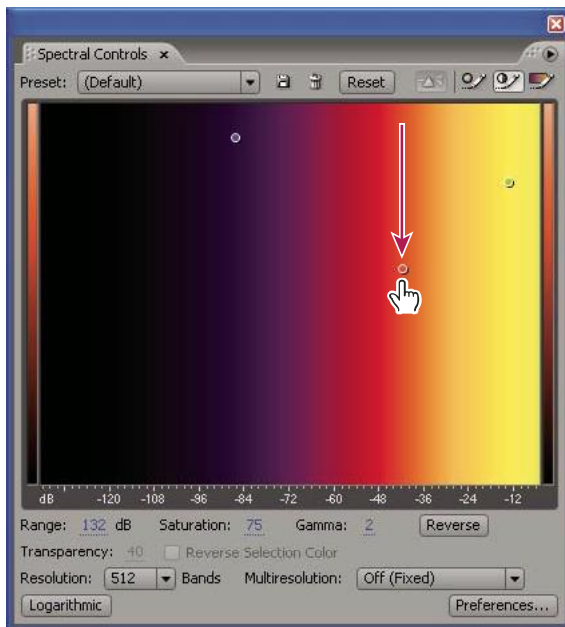
“Techniques for restoring audio” on page 141

Adjusting spectral displays

The Spectral Controls panel (Window > Spectral Controls) gives you precise control over how data is displayed in Spectral Frequency Display, Spectral Pan Display, and Spectral Phase Display. By changing the color and display settings, you can enhance different types of detail and better isolate artifacts.

The Spectral Controls panel lets you modify your display settings dynamically, so you can see your changes immediately in any of the spectral views. You can use one of the display presets or fine-tune the settings to your own specifications. Customized settings can be saved as presets.

In the Spectral Controls panel, colors on the left side of the color spectrum represent low amplitudes, and those on the right side represent high amplitudes. Color points let you precisely adjust the color range. You can also adjust the color saturation and gamma.



Adjusting width of color band in Spectral Controls panel

See also





“About Spectral Frequency Display” on page 66


“About Spectral Pan Display” on page 66

“About Spectral Phase Display” on page 67

“Techniques for restoring audio” on page 141

Customize a spectral display

- 1 In Edit View, choose View > Spectral Frequency Display, Spectral Pan Display, or Spectral Phase Display.
 - 2 Choose Window > Spectral Controls to open the Spectral Controls panel.
 - 3 Choose a preset from the pop-up menu at the top of the Spectral Controls panel.
 - 4 To further adjust the colors, do any of the following:
 - To pan the color spectrum left or right, click anywhere in the color spectrum and drag to the left or right. The colors displayed depend on the Gamma value you specify.
 - To adjust the color and amplitude relationships, drag a color point horizontally.
-  *To constrain movement of a color point either horizontally or vertically, hold the Shift key and then drag the color point.*
- To edit the hue of each color band in the color spectrum, click the Edit Hue button  and then drag a color point vertically.
 - To adjust the brightness of each color band, click the Edit Brightness button  and then drag a color point vertically. Drag up for brighter colors or down for more muted colors.
 - To adjust the color width of bands in the color spectrum, click the Edit Color Width button  and then drag a color point vertically. Drag up to narrow the color width or down to widen it.

- To add color points, right-click the color spectrum and choose Add New Color Point. To delete a color point, right-click it and choose Delete Color Point.
 - (Spectral Pan Display and Spectral Phase Display only) To tint the display based on the actual frequencies in the audio signal, click the Tint Using Spectral Prism button .
 - To restore color settings to the defaults, click Reset.
- 5 (Optional) Set additional color and display options at the bottom of the Spectral Controls panel.
- 6 (Optional) Save the customized settings as a preset for future use.

Spectral Controls panel options

You can set the following options at the bottom of the Spectral Controls panel:

Range Changes the amplitude over which the colors are displayed. Lower the range to see more of the colors at the left end of the spectrum. Increasing the range may help you see more detail in quieter areas of the audio.



You can also adjust the range by dragging the horizontal ruler. Shift-drag to adjust color points relative to the range.

Saturation Changes the overall strength or purity of the colors. Increasing the saturation can bring out more detail.

Gamma Adjusts the steepness of the brightness curve. Higher gamma values produce a steeper brightness curve, which can help add contrast to areas of interest. Lower gamma values produce brighter color in the lower decibel ranges.



You can also adjust the gamma value by Shift-dragging in the color spectrum. Ctrl-drag to adjust the gamma curve's center point.

Reverse Flips the spectral display so that colors on either end of the spectrum switch positions.

Transparency Determines the visual transparency of selected ranges.

Reverse Selection Color Completely inverts selected ranges relative to the background.

Resolution Specifies the number of vertical bands used in drawing frequencies. The higher the number, the more accurate the display, but the longer it will take to render the display.



Multiresolution (Spectral Frequency Display only) Specifies the trade-off between frequency detail and time detail. Multiresolution settings work best if you use a high Resolution value with Logarithmic mode. Choose Off (Fixed) for fixed resolution over the entire frequency range, or choose Time Important or Frequency Important to favor time or frequency detail, respectively. General Purpose balances between Time Important and Frequency Important.

Logarithmic (Spectral Frequency Display only) Toggles the display between logarithmic and linear display.

Quality (Spectral Pan Display and Spectral Phase Display only) Specifies the display quality. Higher quality displays are more accurate, but they take longer to render.

Save or delete a Spectral Controls preset

❖ In the Spectral Controls panel, do one of the following:

- To save custom settings as a preset, click the Save button . Enter a name, select Save Colors Only to save only settings in the color graph, or deselect this option to save all settings. Then click OK.
- To delete a preset, choose it from the Preset menu, and click the Delete button .

Change the vertical scale

You can change the scale of the vertical ruler in Waveform Display and Spectral Frequency Display.

1 Choose View > Vertical Scale Format.

2 Do one of the following:

- In Spectral Frequency Display, choose a Linear or Logarithmic frequency scale.
- In Waveform Display, choose one of the following amplitude scales:

Sample Values Indicates amplitude on a scale that shows the range of data values supported by the current bit depth.

Normalized Values Indicates amplitude on a normalized scale that ranges from -1 to 1.

Percentage Indicates amplitude on a percentage scale that ranges from -100% to 100%.

Decibels Indicates amplitude on a decibel scale that ranges from -Infinity to zero dBFS.



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

See also


“About Waveform Display” on page 65

“About Spectral Frequency Display” on page 66

View the top and tail of an audio file

To create smooth loops, the end of a file should seamlessly transition to the beginning. The Top/Tail view helps you focus on these key areas, while also displaying the entire file for reference.

At the top of the Main panel, time displays show how much audio is visible for the top and tail. By default, Adobe Audition displays 10 seconds in both views. To change the default, choose Edit > Preferences, click the Display tab, and enter an Initial Viewing Range. (If a file is shorter than twice this setting, no more than 45% of total length appears in each view.)






Top/Tail view


A. Top view B. Time visible for top and tail C. Tail view D. Complete audio file with scroll bar above

1 Choose View > Top/Tail Views.

2 To zoom in or out on the top and tail, do any of the following:

- Click the Zoom In  or Zoom Out  button in the upper-left or upper-right corner of the Main panel.
- Right-click and drag in the ruler below the top or tail.
- Position the mouse over the top or tail, and roll the scroll wheel.

 *By default, the top and tail zoom together. To zoom them independently, choose Edit > Preferences, click the Display tab, and deselect Link Top/Tail Zoom Levels.*

3 Start playback to play the top and tail in a continuous loop. (To change this default playback method, right-click the Play button  in the Transport panel, and choose a different option.)

See also


“Playing audio” on page 60

“Zooming audio” on page 20

“Navigate with a scroll bar or ruler” on page 21

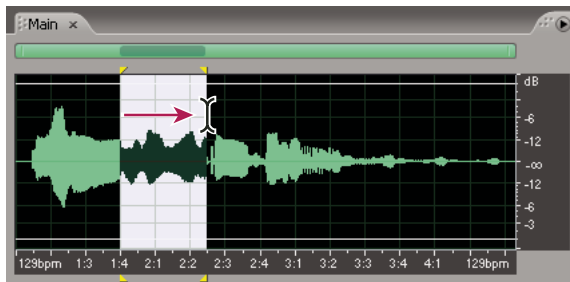
Selecting audio

Select time ranges

1 In the toolbar, select the Time Selection tool .

2 Do any of the following:

- To select a range, drag in the Main panel.



Dragging to select time ranges


- 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 Preferences dialog box. (See “General preferences” on page 33.)*




Select spectral ranges

When working in a spectral display, you can use the Marquee Selection, Lasso Selection, or Effects Paintbrush tool to select audio data within specific spectral ranges. The Marquee Selection tool lets you select a rectangular area. The Lasso Selection and Effects Paintbrush tools (available only in Spectral Frequency Display) let you make free-form selections. All three tools allow for detailed editing and processing, including incredible flexibility in audio restoration work. For example, if you find audio artifacts, you can select and edit just the affected frequencies, producing superior results with faster processing.

The Effects Paintbrush tool creates unique selections that determine the intensity of applied effects. To adjust intensity, either layer brush strokes or change the Opacity setting in the toolbar. The more opaque the white, selected area is, the more intense applied effects will be.



Types of spectral selections
A. Marquee B. Lasso C. Effects Paintbrush




1 In a spectral display, select the Marquee Selection tool , the Lasso Selection tool , or the Effects Paintbrush tool  in the toolbar.

2 Drag in the Main panel to select the desired audio data.

Note: When you make a selection in a stereo waveform, the selection is applied to both channels by default. To select audio data in just one channel, choose *Edit > Edit Channel*, and then choose *Edit Left Channel* or *Edit Right Channel*.

3 To adjust the selection, do any of the following:

- To move the selection, position the pointer in the selection, and drag it to the desired location.
- To resize the selection, position the pointer on the corner or edge of the selection, and drag it to the desired size. (For paintbrush selections, you can also adjust the Brush Size setting in the toolbar.)
- To add to a lasso or paintbrush selection, Shift-drag. To subtract from the selection, Alt-drag.
- To determine the intensity of effects applied to paintbrush selections, adjust the Opacity setting in the toolbar.
- To select audio above and below marquee selections, select Exclude Selection in the toolbar. (For example, if you select center-panned audio, this option excludes that audio and instead selects left- and right-panned audio in the same time range.)

 By default, Adobe Audition plays all audio in the same time range as a spectral selection. To hear only the selection, right-click the standard  or circled  Play button, and select *Play Spectral Selection*.

See also

“About Spectral Frequency Display” on page 66


“About Spectral Pan Display” on page 66

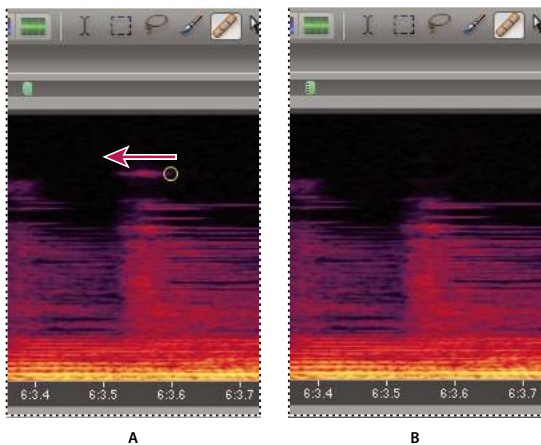
“About Spectral Phase Display” on page 67

“Techniques for restoring audio” on page 141

Select artifacts and repair them automatically

For the quickest repair of small, individual audio artifacts like isolated clicks or pops, use the Spot Healing Brush. When you select audio with this tool, it automatically applies the Favorites > Auto Heal command.

- 1 In Spectral Frequency Display, select the Spot Healing Brush  in the toolbar.
- 2 In the toolbar, adjust the Brush Size setting to change the pixel diameter.
- 3 In the Main panel, click an audio artifact or drag across it.



Instantly removing an artifact with the Spot Healing Brush
A. Before B. After

See also

“Adjusting spectral displays” on page 68

“Techniques for restoring audio” on page 141




Select all of a waveform

- ❖ Do either of the following:
 - To select the visible range of a waveform, double-click in the Main panel.
 - To select all of a waveform, choose Edit > Select Entire Wave, or triple-click in the Main panel.

Specify which channel you want 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.

- ❖ Do any of the following:
 - To select a range in only one channel, either drag near the top of the left (upper) channel or near the bottom of the right (lower) channel. The cursor displays an *L* icon or an *R* icon to indicate the channel in which you are selecting audio.

- 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 shortcut bar. (See “Display the shortcut bar” on page 27.)

Adjust a selection 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.

- ❖ 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.

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.

Select audio between beats

For some editing tasks, such as constructing drum loops and similar musical phrases, you need to select audio between beats. Though you can often identify beats by looking for peaks in a waveform, the Find Beats command makes the task much quicker. After you find beats with this command, Adobe Audition saves them as Beat Markers, making it easy to locate beats again.

- 1 Click in the Main panel to the left of the first beat you want to find.
- 2 Choose Edit > Find Beats > Find Next Beat (Left Side). The cursor moves to the beginning of the next beat.
- 3 Choose Edit > Find Beats > Find Next Beat (Right Side) to select from the current cursor position to the next beat.
- 4 If you want to select more than one beat, choose Edit > Find Beats > Find Next Beat (Right Side) again. Each time you choose this command, Adobe Audition adds the next beat to the selection.

See also

“Marker types” on page 81

“Set markers automatically” on page 84

Change Find Beats settings

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



To find beats in material that has fast transient attacks, such as drums, specify a quick Rise Time and a high Decibel Rise. For material with softer attacks, such as bass, the Rise Time can be slightly slower relative to Decibel Rise.

Snap to markers, rulers, frames, and zero crossings

Snapping causes selection boundaries, as well as the start-time indicator, to move to items such as markers, 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:

Snap To Markers Snaps to a marker point. For information on defining markers, see “About markers” on page 80.

Snap To Ruler (Coarse) Snaps 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) Snaps 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 Snaps to the nearest place where audio crosses the center line (the zero amplitude point).

Snap To Frames (Always) Snaps to a frame boundary if 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.

See also

“Snap to clip and loop endpoints” on page 193

Copying, cutting, pasting, and deleting audio

Choose 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.

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



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

Copy or cut audio data

1 In Waveform Display of Edit View, select the audio data you want to copy or cut. Or, to copy or cut the entire waveform, deselect all audio data.

2 Choose one of the following:


- Edit > Copy, to copy audio data to the active clipboard.
- Edit > Copy To New, to copy and paste the audio data into a newly created file in one step.
- Edit > Cut, to remove audio data from the current waveform and copy it to the active clipboard.

Paste audio data

The Paste command puts any audio data in the active clipboard into the current waveform. The Paste To New command lets you create a new file and insert audio data from the active clipboard.


❖ Do either of the following:

- To paste audio data into the current file, place the cursor where you want to insert the audio data or select the audio data you want to replace. Then choose Edit > Paste. 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.
- To paste audio data into a new file, choose Edit > Paste To New. The new file automatically inherits the properties (sample rate, sample frequency, and so on) from the original clipboard material.

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


Mix 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.*

- 1** In the Main panel, 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.
- 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.

Delete or trim 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.


❖ Do one of the following:

- To delete audio data, select the audio data you want to delete. Then choose Edit > Delete Selection.
- To trim audio data, select the audio data you want to keep. Then choose Edit > Trim.


Visually fading and changing amplitude

About on-clip fade and gain controls

Though various effects can change amplitude or produce fades, on-clip fade and gain controls make the task quick and intuitive. As you drag these controls in the Main panel, a visual preview helps you precisely adjust audio.




*On-clip controls in the Main panel
 A. Fade controls B. Gain control*

 *If you find the on-clip controls visually distracting, deselect On-clip UI in the View menu.*

Visually fade in or out

Adobe Audition offers three types of visual, on-clip fades:


- Linear fades produce an even volume change that works well for much material. If this fade sounds too abrupt, however, try one of the other options.
- Logarithmic fades smoothly change volume slowly and then rapidly, or vice versa.
- Cosine fades are shaped like an S-curve, changing volume slowly at first, rapidly through the bulk of the fade, and slowly at the finish.



*Fade types
 A. Linear B. Logarithmic C. Cosine*

❖ In the upper left or right of the waveform, drag the Fade In  or Fade Out  handle inward, and do any of the following:

- For a linear fade, drag perfectly horizontally.
- For a logarithmic fade, drag up or down.
- For a cosine (S-curve) fade, press Ctrl.

 *To create cosine fades by default and press Ctrl to create linear or logarithmic fades, choose Edit > Preferences, click the General tab, and change the Default Fade setting.*

See also

“Amplify/Fade effect (Edit View only)” on page 112

“Envelope effect (Edit View only)” on page 116

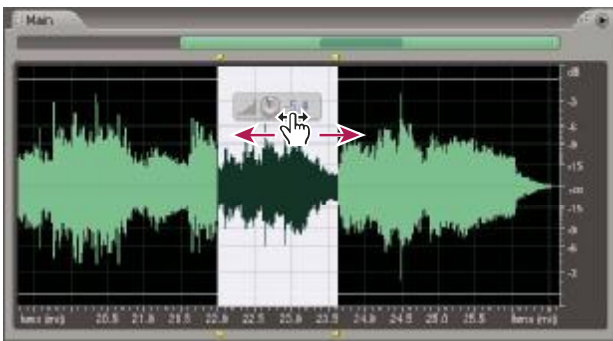
“General preferences” on page 33

Visually raise or lower amplitude

1 In the Main panel, select the audio you want to adjust. (To select an entire file, triple-click.)

2 In the on-clip gain control above the selected audio, drag the blue numbers.

The blue numbers indicate how new amplitude compares with existing amplitude. When you release the mouse button, the numbers return to 0 dB, so you can make further adjustments.



Changing the volume of selected area



To adjust entire files with the on-clip gain control, without creating a selection, choose *Edit > Preferences*, click the *Display* tab, and select *Show On-Clip Gain Control When There Is No Selection Range*.

See also

“Amplify effect” on page 112

“Display preferences” on page 36


Working with markers

About markers

Markers are locations that you define in a waveform. Markers make it easy to navigate within a waveform to make a selection, perform edits, or play back audio.

In Adobe Audition, a marker can be either a *point* or a *range*. A point refers to a specific time position within a waveform (for instance, 1:08.566 from the start of the file). 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). You can drag start and end markers for a range to different times.

Markers have triangular handles that appear at the top and bottom of the Main panel. You use marker handles to select and adjust markers. You can also right-click a marker handle to access commands for working with markers.



Examples of markers

A. Marker handle B. Marker point C. Marker range D. Nonsplit marker range

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

Marker types

Adobe Audition provides five marker types. All five can be ranges as well as points. You can change marker types by clicking *Edit Marker Info* in the *Markers* panel and choosing a marker type from the *Type* menu. (Alternatively, right-click the marker handle, and choose a marker type from the context menu.)

Cue markers Mark important sections of a waveform for later reference (for example, to identify an editing point).

Beat markers Similar to Basic markers, but also identify musical beats for loops, providing very accurate beat mapping. (See “About loops” on page 221.)

Track markers Indicate tracks for an audio compact disc. Use these markers only for burning CDs. (See “Insert CD tracks” on page 264.)

Index markers Set markers within a CD track. (Some CD players can navigate to index markers in addition to track markers.)



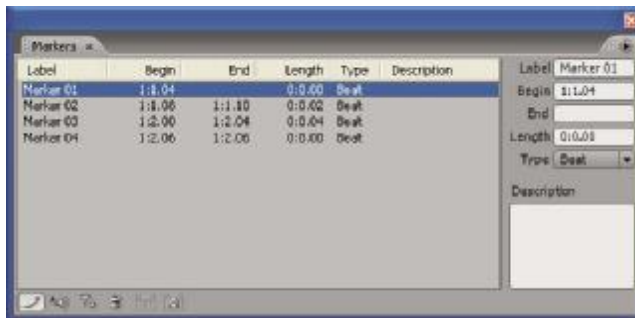
For the smoothest CD-creation workflow, specify ranges for Track markers and points for Index markers.

BWF-J markers Specify marker points for Broadcast Wave files for the Japanese broadcasting environment.

Note: To specify a marker as BWF-J, you must first name it with one of the following labels: *BC\$START*, *BC\$STANDBY*, *BC\$CM*, *BC\$END*, *BC\$STOP*, *BC\$STOP*, *BC\$FILE*, *BC\$PAUSE*, or *BC\$UTL [*]*, where * represents the number 1, 2, 3, or 4.

Add, select, and rename markers

Though you can add markers directly in the *Main* panel, you use the *Markers* panel (*Windows > Marker List*) to define and select markers.




Markers panel

Add a marker



1 Do one of the following:

- Play audio.
- Place the cursor exactly where you want the marker point to be.
- Select the audio data you want to define as a marker range.

2 Either press the F8 key, or click the Add Marker button  in the Markers panel.

Select a marker

- Double-click a marker handle in the Main panel.
- Double-click a marker in the Markers panel.
- To select adjacent markers, click the first marker you want to select in the Markers panel, and then Shift-click the last.
- To select nonadjacent markers, Ctrl-click them in the Markers panel.

 To automatically play markers when you select them in the list, click the Auto Play button . This option helps you quickly identify marker locations.

Rename a marker and add a description

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

1 Select a marker.

2 Click Edit Marker Info in the Markers panel.

3 Do any of the following:

- Enter a new name in the Label text box.
- Enter a description in the Description text box.

Adjust, merge, convert, or delete markers

After creating markers, you can fine-tune them to best address the needs of an audio project.

Reposition point or range markers


- In the Main panel, drag the marker handle (for point markers) or the red start handle (for range markers) to a new location.

- In the Markers panel, select the marker, and enter a new value in the Begin text box.

Change the duration of a range marker


- In the Main panel, drag the blue end handle to a new location.
- In the Markers panel, select the marker, and enter a new value in the End or Length text box.

Merge individual markers

- ❖ Open the Markers panel, select the markers you want to merge, and click the Merge button .

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

Merge two marker ranges

- ❖ Open the Markers panel, select the two ranges, and click the Merge button .

Convert a point marker to a range marker

- ❖ Right-click the marker handle and choose Make Range.


The marker handle splits into two handles.

Convert a range marker to a point marker

- ❖ Right-click a marker handle and choose Make Point.

The two parts of the range marker handle merge into a single handle, with the start time of the range becoming the time for the point marker.

Delete markers

- Select one or more markers, and click the Delete button  in the Markers panel.
- Right-click the marker handle in the Main panel, and choose Delete.

Add silence between markers


- 1 In Edit View, choose Window > Marker List.
- 2 Select one or more markers. At least one of the markers you select must be a range.
- 3 Click Batch in the Markers panel.
- 4 Click Set Amount Of Silence, enter the desired number of seconds in the Add Silence Before and Add Silence After boxes, and click OK.

See also

“Save audio between markers to new files” on page 83

“Create silence” on page 86

Save audio between markers to new files

- 1 In Edit View, choose Window > Marker List.
- 2 Select one or more markers. At least one of the markers you select must be a range.
- 3 Click the Batch process button  in the Markers panel.

4 Click Save To Files, set the following options as desired, and click OK:

Use Marker Label As Filename Uses the marker name as the prefix for the filename.

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 for the output format you choose.

Seq. Start Specifies the number to begin with when adding numbers to the filename prefix.

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

Output Format Specifies the output format for the new files. If the specified format has unique settings, you can click Options to access them.

See also

“Add silence between markers” on page 83

“About audio file formats” on page 244

“Batch processing overview” on page 259

Set markers automatically

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

1 Select the general range in which you want to find phrases or beats.

2 Choose Edit > Auto-Mark, 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 marker list.

Find Phrases And Mark Scans the selected range, marking nonsilent ranges as basic markers in the marker list.

Find Beats And Mark Scans the selected range, marking beats as beat markers in the marker list.

Trim Digital Silence Trims silence from the beginning and ending of a file. 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.

Customize Auto-Mark settings

1 Choose Edit > Auto-Mark > Auto-Mark 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 work well—for example, words or phrases get chopped off—lower the signal level values. Increase the signal level values if not enough silence is removed.

Creating play lists

A *play list* is an arrangement of marker ranges that you can play back in any order and loop a specified number of times. A play list lets you try different versions of an arrangement before you commit to edits. You create play lists in the Play List panel (Window > Play List).

See also

“Working with markers” on page 80

Create a play list

- 1 If the Markers panel isn't visible, click Open Marker List in the Play List panel.
- 2 In the Markers panel, select the marker ranges you want to add to the play list.
- 3 In the Play List panel, click Insert Marker. The range is inserted either before the currently selected play list item or at the end of the list if nothing is selected.

Play items in a play list

- To play the entire list, select the first item, and click Play Markers.
- To play part of the list, select the first item you want to play, and click Play Markers.
- To play a specific item, select it, and click Play Current Item.

Change the order of items in a play list

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

Loop an item in a play list

- ❖ Select an item, and enter a number in the Loops box. Each item can loop a different number of times.

Delete items from a play list

- ❖ Select the items, and click Remove.

Creating and deleting silence

Create silence

Creating silence is useful for inserting pauses and removing nonessential noise from an audio file. 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.

❖ Do either of the following:


- To mute existing audio data, select the desired range of audio data and choose **Effects > Mute**. Unlike deleting or cutting a selection, which splices the surrounding material together, applying the Mute effect leaves the duration of the selection intact, and simply zeros the amplitude within it.
- To insert a new duration of silence, 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. Then choose **Generate > Silence** and 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. Click **OK**. Any audio to the right of the cursor is pushed out in time, thereby lengthening the waveform's duration.

Delete 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.

- 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 Marker List Adds each location where silence is removed to the marker 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

Invert a waveform

The Invert effect inverts audio phase by 180 degrees. Inverting doesn't produce an audible change on an individual waveform, but you can hear a difference when combining waveforms. For example, you might invert pasted audio to better align it with existing audio. Or, you could invert one channel of a stereo file to correct an out-of-phase recording.

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

See also

“Waveform measurements” on page 8

“phase” on page 277

Reverse a waveform

The Reverse effect reverses a waveform from right to left so it plays backwards. Reversing is useful for creating special effects.

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

Generating audio

Generate DTMF signals

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

- 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 (Process).
- 3 Set options as desired, and click OK.

DTMF options

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.

Generate 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.

- 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 options as desired, and click OK.

Noise options

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.

Generate 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.

- 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.

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° , waves will start at the baseline. If Start Phase is set to 90° , 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° 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.

Analyzing phase, frequency, and amplitude

Techniques for analyzing audio

Edit View provides several ways to analyze audio. To compare phase relationships between the left and right stereo channels, use the Phase Analysis panel. To analyze tonal and dynamic range, use the Frequency Analysis panel and the Amplitude Statistics dialog box.

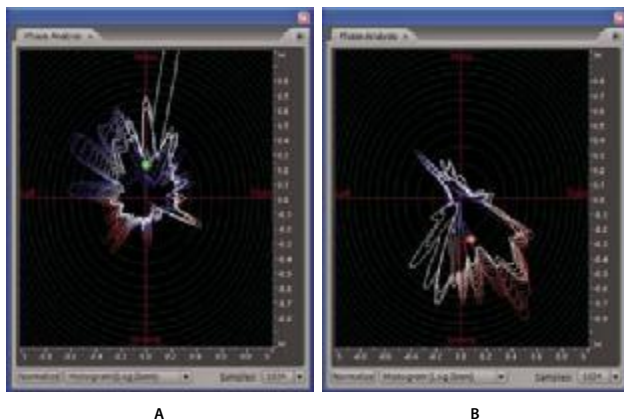
Edit View also provides Spectral Frequency, Spectral Pan, and Spectral Phase displays, which you can use in conjunction with the analysis methods above. (See “Displaying audio in Edit View” on page 65.)

Analyze phase

In Edit View, you can use the Phase Analysis panel to reveal out-of-phase channels for stereo waveforms, which you can correct with the Invert command. Phase analysis also shows stereo separation, which you can adjust using stereo imagery effects. (See “Changing stereo imagery” on page 162.)



To evaluate phase in detail, use the Phase Analysis panel in conjunction with Spectral Phase Display. (See “About Spectral Phase Display” on page 67.)



Phase Analysis panel using Histogram (Log Zoom) view:
A. In-phase audio B. Out-of-phase audio

1 In Edit View, choose Window > Phase Analysis.

2 In Waveform Display of Edit View, click a point, select a range, or start playback.

If you select a range, Adobe Audition analyzes the center point of the range.

3 Use the following display items to interpret the phase information shown in the Phase Analysis panel:

- By default, the graph shows stereo channels on the horizontal axis and phase information on the vertical axis. In-phase audio appears in the top half of the graph; out-of-phase audio appears in the bottom half.



To zoom in on the graph, use the mouse wheel or right-click a ruler and drag the magnifying glass icon. (See “Zoom phase and frequency graphs” on page 95.)

- The colored ball shows the overall stereo and phase location of audio. The ball turns red for out-of-phase audio, green for in-phase audio, and yellow for borderline audio.
- The white outline shows the percentage of samples in each channel that are in- or out-of-phase, reflecting the width of the stereo image.



If you plan to create a mono version of a stereo file (for AM radio, for example), note the width of the stereo image. If it's very broad, the file will sound significantly different in mono.

- The colored tracer lines indicate whether the waveform is more in phase (blue) or out of phase (red).

See also

“Invert a waveform” on page 87

“About Spectral Phase Display” on page 67

Phase Analysis options

Set the following options to customize the Phase Analysis panel:

Normalize Adjusts the phase analysis lines so that they reach the edge of the graph.

Chart Type menu Select from the following options:

- **Left/Right** Displays a Lissajou Plot graph in which 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; and a stereo waveform with wide separation appears as many wavy lines extending in all directions.

- **Mid/Side** Rotates the Left/Right display to the left by 45°. 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** Displays amplitude rather than phase, redrawing the waveform on an axis that rotates in the phase graph.
- **Histogram (Log Zoom)** Displays the default circular histogram.
- **Histogram (Log Full)** Similar to Histogram (Log Zoom), but displays the entire range. This option is useful for viewing pure mono at the edges of the graph.
- **Histogram (Linear)** Similar to Histogram (Log Zoom), but displays the graph on a linear (rather than logarithmic) scale. You may notice larger jumps and variations when using this option because of the linear scale. Click Normalize to constrain the display to the full screen.
- **Phase Wheel** Displays the phase of individual frequencies. Higher frequencies appear toward the outer edge of the 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.

Additional display options for Histogram modes Right-click the graph and select from the following options:

- **Show Tracers** Displays blue and red tracer lines that help you see in- and out-of-phase audio, respectively.
- **Show Ball** Displays a ball that shows the overall stereo and phase location of audio. The ball turns red for out-of-phase audio, green for in-phase audio, and yellow for borderline audio.
- **Modulate By Amplitude** Changes the size of the phase graphic relative to amplitude. When this option is selected, the phase graphic shrinks for quiet audio and gets larger for loud audio.
- **Show Labels** Displays the Mono, Left, Right, and Inverse axis labels.
- **Show Peaks** Displays a yellow line that slowly contracts from the highest recent peak.
- **Show Valleys** Displays a purple line that slowly expands from the lowest recent valley.
- **Static Peaks and Valleys** Displays the absolute peak and valley since the start of audio playback.


Additional display option for Spin mode Right-click the graph and select Edit Spin Rate to change the speed of rotation.

Analyze frequency range

In Edit View, you can use the Frequency Analysis panel to identify problematic frequency bands, which you can then correct with a filter effect. In this panel, the horizontal axis represents frequency (measured in Hz), while the vertical axis represents amplitude (measured in decibels).



To save or print a Frequency Analysis graph, press Alt+PrtScn, and then paste into an image-editing application.




Frequency Analysis panel displaying Advanced options
A. Musical note B. Frequency scale C. Chart format D. Amplitude scale E. FFT window type

- 1 In Edit View, choose Window > Frequency Analysis.
- 2 In the Main panel, click a point, select a range, or start playback.

If you select a range, Adobe Audition analyzes only the center point. To analyze the overall frequency of a selected range, click Scan Selection in the Frequency Analysis panel.

- 3 View frequency and amplitude information in the Frequency Analysis panel:

- The graph displays frequency along the horizontal axis and amplitude on the vertical axis.

 To zoom in on a particular area of the graph, right-click on a ruler and drag the magnifying glass icon. (See “Zoom phase and frequency graphs” on page 95.)

- The musical note of the waveform is displayed in the area above the graph. The numbers indicate the keyboard position and variance from standard tuning. For example, A2 +7 equals the second-lowest A on a keyboard tuned 7% higher than normal.
- The highest frequency of the entire waveform and the maximum amplitude for each channel is displayed in the left status area underneath the graph.
- The overall frequency (and equivalent musical note) at the center point of the selected range is displayed in the right status area underneath the graph.

Frequency Analysis options

Set the following options to customize the Frequency Analysis panel:

Linear View Sets the graph display to a linear horizontal frequency scale when selected or a logarithmic scale when deselected.

Hold buttons Let you 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.

Chart Format menu Changes the graph display. Choose one of the following styles:

- **Lines** Displays amplitude at each frequency with simple lines. The left channel is blue; the right is red.

- **Area (Left On Top)** 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.

FFT Size menu (Advanced option) 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 panel 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 panel.

Note: When you set the FFT size to 8192 or lower, the Frequency Analysis panel 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 Window Type menu (Advanced option) Lets you choose from eight types of FFT windows. Each displays a slightly different kind of frequency graph.

The Triangular option 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 option displays 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 (Advanced option) 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.

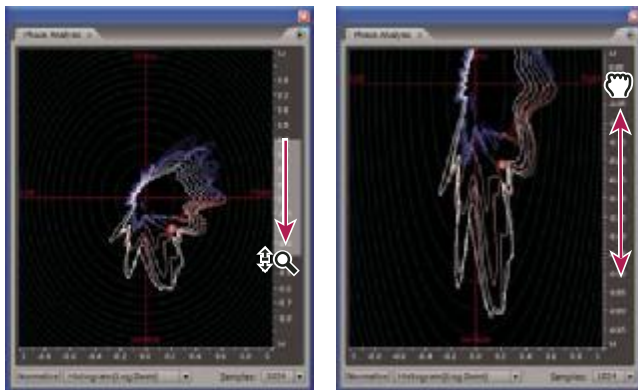
Copy To Clipboard (Advanced option) Copies a text-based frequency report of the current waveform to the Windows clipboard.

Zoom phase and frequency graphs

In the Phase Analysis and Frequency Analysis panels, you can zoom graphs to analyze phase and frequency in more detail.

❖ Do any of the following:

- To zoom in on a graph, right-click and drag the magnifying glass icon in the vertical or horizontal ruler. (In the Phase Analysis panel, you can also click in the panel and use the mouse wheel to zoom.)
- To navigate a magnified graph, left-click and drag the hand icon  in the vertical or horizontal ruler.
- To zoom out on a magnified graph, right-click in the vertical or horizontal ruler, and choose Zoom Out or Zoom Out Full. Zoom Out to return to the previous magnification. (This option is available only in the Frequency Analysis panel.) Zoom Out Full to zoom out completely.




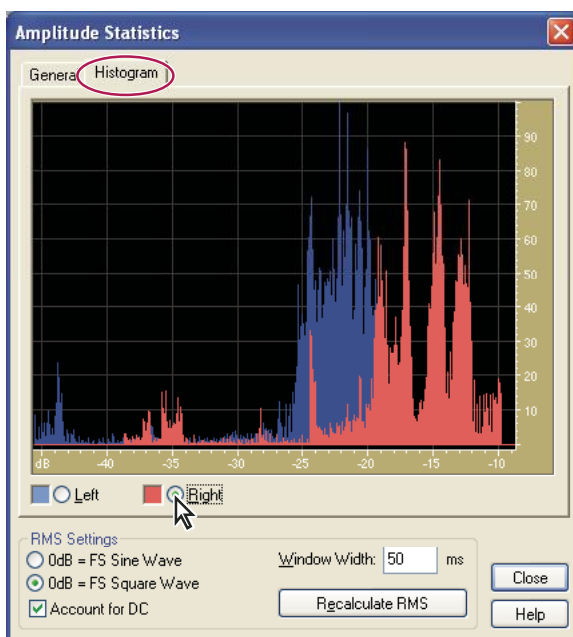
Zooming and navigating a Phase Analysis graph

Analyze amplitude

In Edit View, you can use the Amplitude Statistics dialog box to evaluate a variety of information about audio amplitude.

- 1 In Edit View, select an audio range.
- 2 Choose Window > Amplitude Statistics, and view the following:
 - 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. Select Left or Right to display either the left or right channel in the foreground.

 Use the Histogram tab to identify prevalent amplitudes, and then compress, limit, or normalize them with an amplitude effect.





Amplitude Statistics dialog box, Histogram tab


Amplitude Statistics options


In the Amplitude Statistics dialog box, you can access specific waveform points from the General tab or change RMS settings to recalculate the statistics.

General tab


Minimum Sample Value Shows the sample with the lowest amplitude. Click the arrow button  to place the start-time indicator at that location and close the Amplitude Statistics dialog box.


Maximum Sample Value Shows the sample with the highest amplitude. Click the arrow button  to place the start-time indicator at that location and close the Amplitude Statistics dialog box.

Peak Amplitude Shows the sample with the highest amplitude in decibel form. Click the arrow button  to place the start-time indicator at that location and close the Amplitude Statistics dialog box.

Possibly Clipped Samples Shows the number of samples that could exceed 0 dBFS. Click the arrow button  to place the start-time indicator at the first clipped sample and close the Amplitude Statistics dialog box. (If necessary, choose Window > Amplitude 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 start-time indicator at that location and close the Amplitude Statistics dialog box.

Maximum RMS Power Shows the maximum RMS amplitude. Click the arrow button  to place the start-time indicator at that location and close the Amplitude 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 section

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.

Converting sample types

Preview a file in a different sample rate

The Adjust Sample Rate command lets you preview how an audio file will sound at a different sample rate. (See “Sample rate” on page 10.) This command doesn’t convert the sample rate of the audio file—use the Convert Sample Type command to do that.

- 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 work with any sample rate in Adobe Audition, your sound card may not be capable of playing all rates properly. To determine supported sample rates, consult the documentation for the card.

Convert the sample rate of a file

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 converting the sample rate, keep in mind that most sound cards support only certain sample rates.

- 1 Choose Edit > Convert Sample Type.
- 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.

See also

“Sample rate” on page 10

Convert a waveform 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.

- 1 Choose Edit > Convert Sample Type.
- 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.*

Create a stereo waveform from two mono waveforms


- 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.

Change the bit depth of a file

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.

 *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.*

- 1 Choose Edit > Convert Sample Type.
- 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, Adobe Audition provides dithering options to help reduce noise and distortion. Set the following options as desired, and click OK:

 *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.*

Enable Dithering Enables or disables dithering. If dithering is disabled, 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 OK 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 trade-off 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	6.0 dB	Negligible
Shaped Triangle	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 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.

See also

“Bit depth” on page 11

Use sample rate conversion presets

If you need to make the same conversion on multiple files, you can save time by using a sample rate conversion preset.

- 1 Choose Edit > Convert Sample Type.
- 2 Adjust the settings as desired.
- 3 Click Save As, type a name for the preset, and click OK.

After you create a preset, it appears in the preset list at the top of the dialog box. To apply the preset to a file, select the preset and click OK. If you want to delete a preset, choose the preset from the list, and click Delete.


Recovery and undo

Undo or redo changes


Each time you start Adobe Audition, it keeps track of the edits you perform. 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.

❖ To undo or redo changes, do any of the following:

- To undo a change, choose Edit > Undo [*name of change*]. 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*].
- To repeat the last command, in Edit View, choose Edit > Repeat Last Command. You can repeat most commands; however, there are a few exceptions (such as Delete).


 *To repeat the last command without opening its dialog box, press the F3 key.*

Disable or enable the Undo function

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.

❖ Do one of the following:

- In Edit View, choose Edit > Enable Undo/Redo. A check mark indicates that the Undo function is enabled.
- Choose Edit > Preferences, 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 preferences" on page 34.)

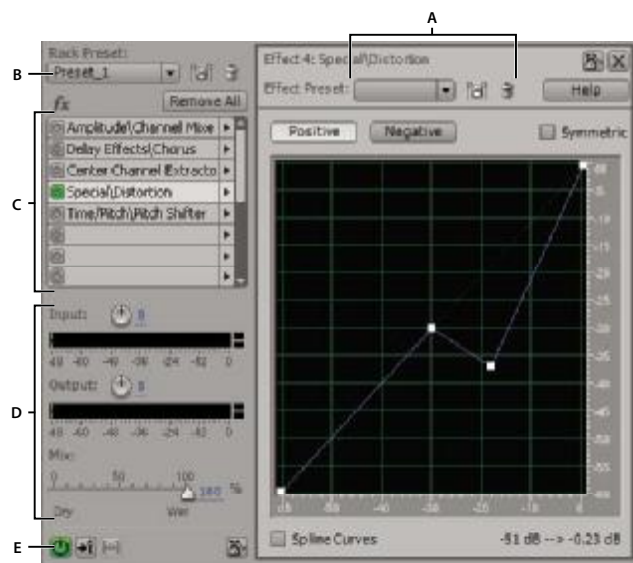
 *If you don't have enough disk space to save the undo information, you can change the Temp folder to a different drive, if available.*

Chapter 7: Applying effects

Shared effects controls

Mastering Rack and Effects Rack overview

Two similar features let you apply groups of effects: the Mastering Rack in Edit View, and the Effects Rack in Multitrack View. In both racks, you can insert, edit, and reorder up to 16 effects, optimize mix levels, and store favorite presets.



Controls shared by the Mastering Rack and Effects Rack
 A. Effect Preset controls B. Rack Preset controls C. Effect slots D. Level controls E. Main Power button

See also


“Apply groups of effects in Edit View” on page 106

“Apply effects in Multitrack View” on page 107

“Comparing Edit View and Multitrack View” on page 20

Controls unique to the Mastering Rack and Edit View


In Edit View, the Mastering Rack provides Preview and OK buttons that you use to hear and permanently apply effects. If you click OK or Close, the Mastering Rack stores your settings so you can quickly reapply them.



Controls unique to the Mastering Rack and Edit View
 A. Preview button with Preroll/Postroll option B. Buttons that apply or cancel permanent edits

Controls unique to the Effects Rack and Multitrack View

In Multitrack View, the Effects Rack provides Safe During Write, Freeze, and Pre/Post Fader buttons that you use to automate, lock, and route effects. Each track has its own Effects Rack, which is saved with the session. (Dock the Effects Rack to quickly view and edit effect settings for multiple tracks.)




Controls unique to the Effects Rack and Multitrack View
A. Pre/Post Fader button B. Freeze button C. Safe During Write button

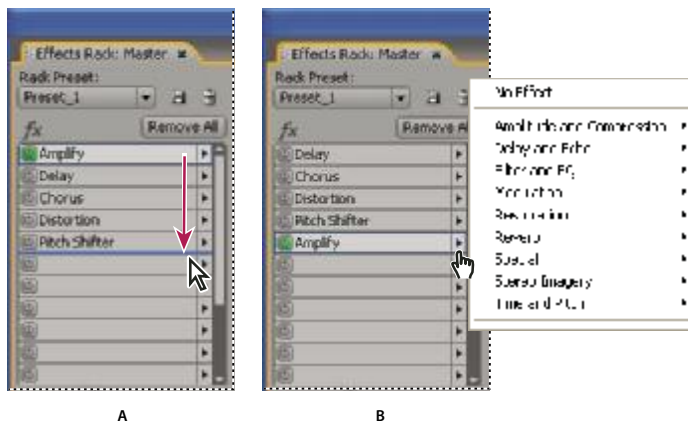
Set input, output, and mix levels in racks

- To optimize volume, adjust Input and Output levels so their meters peak without clipping.
- To change the percentage of processed audio, drag the Mix slider. 100% (Wet) equals fully processed audio; 0% (Dry) equals original, unprocessed audio.


Insert, bypass, reorder, and remove effects in racks


In the Mastering Rack and Effects Rack, you manage groups of effects by using individual effect slots.

 In Multitrack View, the FX section of the Main panel or Mixer provides quick access to slots in the Effects Rack.



Reordering and inserting effects in racks:
A. Reorder by dragging B. Insert with the slot menu

- To insert an effect, choose it from a slot's pop-up menu.
- To bypass an effect, click its Power button .
- To bypass all effects, click the main Power button in the lower left corner of a rack, or the FX Power button in the Main panel or Mixer.

 Bypass effects to quickly compare processed and unprocessed audio.

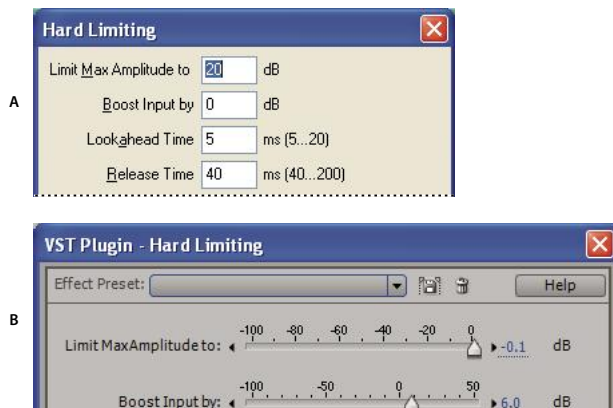
- To remove a single effect, choose No Effect from a slot's pop-up menu.
- To remove all effects, click Remove All. (To access this option in the Main panel or Mixer, choose Effects Rack from a slot's pop-up menu.)
- To reorder effects, drag them to different slots.

💡 *Reordering effects produces different sonic results. (For example, place Reverb prior to Sweeping Phaser, and vice versa.)*

About process effects

Unlike VST effects, process effects can only be applied individually, so they aren't available in the Mastering Rack or Effects Rack. However, some process effects provide more options than similar VST versions. (For example, compare the process and VST versions of the Hard Limiting and Stereo Field Rotate effects.)



You can identify process effects by the word *process* in menu commands and by visual appearance. Dialog boxes for process effects feature a different color scheme and option arrangement than dialog boxes for VST effects.



*Process and VST effects differ in appearance.
A. Process effect B. VST effect*

Use effect presets

Many effects provide presets that let you store and recall favorite settings. In addition to effect-specific presets, the Mastering Rack and Effects Rack provide rack presets that store groups of effects and settings.

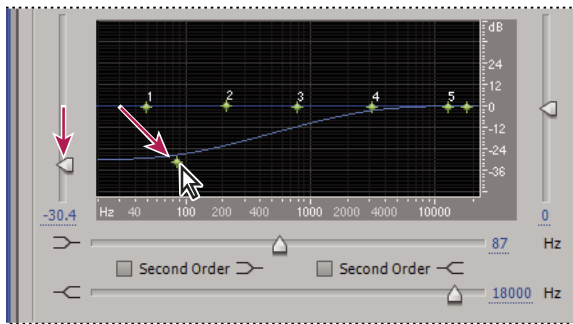
- To apply a preset, choose it from the Preset menu. (In dialog boxes with a list rather than a menu, click the preset name.)
- To save current settings as a preset, click the Save button  (labeled “Add” in some dialog boxes).
- To delete a preset, select it, and click the Delete button  (labeled “Del” in some dialog boxes).

💡 *To modify an existing preset, apply it, adjust settings as desired, and then save a new preset with the same name.*

Control effects settings with graphs

Many Adobe Audition effects provide graphs where you can adjust parameters. By adding and moving control points on the graph, you can precisely tailor effect settings.

Graph control points function together with related sliders. If you move or disable the slider for an effects setting, the related graph control follows suit.



Moving a control point moves the related slider, and vice versa.

- To move a point on a graph, drag it to a new location. (When the pointer is over a control point, the pointer changes from an arrow to a hand.)



Note: The following techniques don't apply to the Full Reverb, Parametric Equalizer, and Track EQ graphs. (See "Full Reverb effect" on page 153, "Parametric Equalizer effect" on page 132, and "Equalizing tracks" on page 190.)

- To add a control point to a graph, click in the grid at the location where you want to place the point.
- To enter numeric values for a control point, right-click the point to bring up the edit box, or double-click the graph's curve.
- To remove a point from a graph, drag it off the graph.
- To return a graph to its default state, click Reset (labeled "Flat" in some dialog boxes).

About spline curves for graphs

By default, graphs display straight lines between control points. However, some graphs provide a Splines or Spline Curves option that creates a curve between control points for smoother transitions.

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




Graph with straight lines compared to graph with spline curves

Group effects in the Effects panel

The Effects panel lists all available audio effects for easy access. You can change the grouping of effects to best meet your needs.

In addition to Adobe Audition effects, you'll find any third-party DirectX plug-ins on your system, plus VST plug-ins in specified folders. (See "Enable VST effects" on page 110.)



Effects panel

❖ Click the buttons at the bottom of the Effects panel:

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 Multitrack Effects are grouped together.


Deselect Group Real-Time Effects To return to the previous view.

Applying effects in Edit View


Apply groups of effects in Edit View

In Edit View, the Mastering Rack lets you apply groups of effects. (It doesn't provide process effects such as Noise Reduction, which must be applied individually.)

- 1 Select the desired audio with the Time, Marquee, or Lasso Selection tool. (Double-click to select the entire file.)
- 2 Choose Effects > Mastering Rack.
- 3 In the FX list, choose effects for up to 16 slots.
- 4 Click the Preview button , and then edit, mix, and reorder effects as needed.

 To compare processed audio to original audio, select and deselect the main Power button in the lower left corner of the rack, or the Power buttons for individual effects.

- 5 When you're ready to apply the changes, click OK.

 To store settings, save a rack preset. (See “Use effect presets” on page 104.)


See also

“Mastering Rack and Effects Rack overview” on page 102

“About process effects” on page 104

Apply individual effects in Edit View

- 1 In the Main panel, select the desired audio.
- 2 From any submenu in the Effects menu, choose an effect.
- 3 Click the Preview button, and then edit settings as needed.

 As you edit settings, watch the Levels panel to optimize amplitude.

- 4 To compare original audio to processed audio, select and deselect the Bypass option.
- 5 When you're ready to apply the changes, click OK.

See also

“Use effect presets” on page 104

“Control effects settings with graphs” on page 104

“About process effects” on page 104

Add preroll and postroll to effects previews

In Edit View, you can add preroll and postroll to effects previews, revealing how processed audio transitions into surrounding audio. This option is especially helpful when previewing effects applied to small selections.

❖ Do either of the following:

- Choose Effects > Enable Preroll And Postroll Preview.
- In an effects dialog box or the Mastering Rack, select Preroll/Postroll Preview.

See also

“Play selections with preroll and postroll” on page 61

Applying effects in Multitrack View

Apply effects in Multitrack View

In Multitrack View, you can apply up to 16 effects to each audio and bus track and adjust them while a mix plays. You can insert, reorder, and remove effects in the Mixer, Main panel, or Effects Rack. Only in the Effects Rack, however, can you edit effects and save favorite settings as presets, which you can apply to multiple tracks.


In Multitrack View, effects are *nondestructive*, so you can change them at any time. To readapt a session for different projects, for example, simply reopen it and change effects to create new sonic textures.



Revealing effect slots in the Main panel


1 To apply an effect, do one of the following:

- Drag it from the Effects panel to a track.
- Choose it from an effect slot menu in the Main panel or Mixer.

 *To reveal effect slots in the Main panel, click the FX button **fx** in the upper left corner, and then zoom in vertically.*

2 In the Effects Rack, choose effects for up to 16 slots in the FX list.

3 Press the spacebar to play the session, and then edit, mix, reorder, or remove effects as needed.

 *To change effect settings over time, use track envelopes. (See “Automating track settings” on page 204.)*

See also

“Mastering Rack and Effects Rack overview” on page 102

Edit settings for previously applied effects


To edit individual settings for effects you previously applied to a track, you must reopen the Effects Rack.


❖ Do either of the following:

- In the Main panel or Mixer, double-click an effects slot, or choose Effects Rack from a slot’s pop-up menu.
- Select a track, and then choose Window > Effects Rack.

Freeze effects and improve performance

After you apply effects to a track in Multitrack View, you can freeze them to conserve processing power, improving system performance for complex mixes. When a track is frozen, you can't edit effects, clips, or envelopes it contains. However, you can quickly defrost a track if you need to change it. (Although freezing tracks takes a small amount of processing time, defrosting tracks is instantaneous.)

- ❖ In the Main panel, Effects Rack, or Mixer, click the Freeze button .






Identifying a frozen track in the Main panel:


A. Unavailable FX Power and Pre-Fader/Post-Fader buttons B. Freeze Track button and clips colored blue

Insert effects before or after sends and EQ

On each track, you can insert effects either pre- or post-fader. Pre-fader effects process audio *before* sends and EQ. Post-fader effects process audio *after* sends and EQ. For most mixes, the default, pre-fader setting works well. The post-fader setting offers signal-routing flexibility for particularly complex mixes.

- ❖ In the FX section of the Main panel or Mixer, click the Pre-Fader/Post-Fader button to insert effects either before sends and EQ , or after .

 If you're editing effect settings in the Effects Rack, click the Pre-Fader/Post-Fader button in the lower-left corner.



Pre- and post-fader effect and send routing for each track:

A. Input gain B. EQ C. Volume D. Mute E. Send F. Effects Rack

See also

“Routing audio to buses, sends, and the Master track” on page 186

“Set up a send” on page 188

Adding third-party plug-ins

About third-party plug-ins

Third-party DirectX and VST plug-ins let you extend the already powerful effects provided with Adobe Audition. Applying plug-in effects is identical to applying built-in effects. For information about plug-in features, consult the documentation provided by the plug-in manufacturer.

To access audio plug-ins in Adobe Audition, you must either enable all of them (DirectX effects) or enable selected plug-ins (VST effects). Be aware that enabling large numbers of plug-ins can cause Adobe Audition to start slowly.

Note: If a third-party effect is incompatible, Adobe Audition adds it to an Unsupported submenu in effects menus.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Automating track settings” on page 204

Enable DirectX effects

1 In Edit View, choose Effects > Enable DirectX Effects.

2 Click Yes to refresh the effects list.

Note: If effects are being used in a multitrack session, close the session.

Adobe Audition scans your system for DirectX plug-ins. After the plug-ins are activated, the Enable DirectX Effects option is removed from the menu.

Enable VST effects

The Add/Remove VST Directory dialog box lists the folders that Adobe Audition scans for VST plug-ins. By default, all plug-ins in specified VST folders are disabled. To optimize performance, enable only the plug-ins you plan to use in Adobe Audition.



If Adobe Premiere® Pro is installed first, Adobe Audition automatically scans for the VST plug-ins provided with Adobe Premiere Pro.


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

Note: If effects are being used in a multitrack session, close the session.

2 Do either of the following:

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

- 3 Click OK, and then click Yes to open the VST Plug-in Manager.
- 4 Select the plug-ins you want to access in Adobe Audition, and then click OK.

 *To directly access the plug-in manager, choose Effects > VST Plug-in Manager. Though you can select VST instruments here, they are unsupported for audio processing. To enable such plug-ins for MIDI sequencing, use the plug-in manager found in the Sequencer panel of Multitrack View.*

See also

“Enable VSTi plug-ins” on page 214

Chapter 8: Effects reference

Amplitude and compression effects

Amplify effect

The Amplitude And Compression > Amplify effect boosts or attenuates an audio signal. Because the effect is VST-based, you can combine it with other effects in the Mastering Rack or Effects Rack.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Amplify options

Left Gain, Right Gain sliders Boost or attenuate the left and right channels.

Link Left And Right Moves the Left Gain and Right Gain sliders together.

Amplify/Fade effect (Edit View only)

The Amplitude And Compression > Amplify/Fade effect produces either constant amplification changes (such as fixed boosts) or precise fades.



Although the Amplify/Fade effect isn't available in Multitrack View, you can use automation lanes to accomplish the same task. (See “Automating track settings” on page 204.)

See also

“About process effects” on page 104

“Apply individual effects in Edit View” on page 107

Amplify/Fade 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.


Dynamics Processing effect

The Amplitude And Compression > Dynamics Processing effect can be used as a compressor, limiter, or expander. As a compressor and limiter, this effect reduces dynamic range, producing consistent volume levels. As an expander, it increases dynamic range by reducing the level of low-level signals. (With extreme expander settings, you can totally eliminate noise that falls below a specific amplitude 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.



Use the Broadcast Limiter preset to simulate the processed sound of a contemporary FM radio station.



Dynamics Processing graph (Classic SoftKnee preset)

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Control effects settings with graphs” on page 104

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Graphic options

Click **Graphic** to visually manipulate dynamics.

Graph 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).

Reset 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 options

Click Traditional to 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 options

Click Attack/Release to determine how quickly signals that cross the threshold are processed.

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** 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** 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 mode** 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.


Noise Gating Completely silences signals that fall below the specified threshold.

Band Limiting options

Click Band Limiting to limit dynamics manipulation to a specified range.


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 half of the current sample rate. (See "Nyquist frequency" on page 277.)

 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).

Envelope effect (Edit View only)

The Amplitude And Compression > 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 automation lanes to accomplish the same task. (See "Automating track settings" on page 204.)



Envelope graph (Bell Curve preset)

See also

"About process effects" on page 104

"Apply individual effects in Edit View" on page 107

"Control effects settings with graphs" on page 104

"Use effect presets" on page 104

"Add preroll and postroll to effects previews" on page 107

Envelope 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.

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.

Hard Limiting VST effect

Like the similar process effect, the Amplitude And Compression > Hard Limiting VST effect drastically attenuates audio that rises above a specified threshold. However, you can combine the VST version with other effects in the Mastering Rack and Effects Rack. In Multitrack View, you can also vary the effect over time by using automation lanes. (See “Automating track settings” on page 204.)

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Hard Limiting VST 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 Pre-amplifies audio before you limit it, making a selection louder without clipping it. As you increase this level, compression increases. Try extreme settings to achieve the loud, high-impact audio heard in radio and TV commercials.

Lookahead Time Sets the amount of time (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 (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.

Hard Limiting process effect (Edit View only)

The Amplitude And Compression > Hard Limiting effect process 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).

See also

“About process effects” on page 104

“Apply individual effects in Edit View” on page 107

Hard Limiting process 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, making a selection louder without clipping it. As you increase this level, compression increases. Try extreme settings to achieve the loud, high-impact audio heard in radio and 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.


Multiband Compressor effect

The Amplitude And Compression > Multiband Compressor effect lets you independently compress four different frequency bands. Because each band typically contains unique dynamic content, multiband compression is a particularly powerful tool for audio mastering.

Controls in the Multiband Compressor let you precisely define crossover frequencies and apply band-specific compression settings. Click Solo buttons to preview bands in isolation, or Bypass buttons to pass bands through without processing. After you fine-tune individual bands, select Link Band Controls to adjust them globally, and then optimize overall volume with the Output Gain slider and Limiter settings.



To change compression settings over time, use automation lanes in Multitrack View. (See “Automating track settings” on page 204.)




Adjusting a crossover frequency in the Multiband Compressor
A. Frequency bands B. Crossover markers C. Bypassed band (no processing applied) D. Amplitude scale E. Frequency scale

See also

- “Mastering effect” on page 161
- “Applying effects in Edit View” on page 106
- “Applying effects in Multitrack View” on page 107
- “Use effect presets” on page 104
- “Add preroll and postroll to effects previews” on page 107

Multiband Compressor options


Crossover Sets the crossover frequencies, which determine the width of each band. Either enter specific Low, Midrange, and High frequencies, or drag the crossover markers above the graph.



Band-specific controls in the Multiband Compressor
A. Solo B. Bypass C. Threshold slider D. Input Level meters E. Gain Reduction meters

Solo buttons Let you hear specific frequency bands. Enable one Solo button at a time to hear bands in isolation, or enable multiple buttons to hear two or more bands together.

Bypass buttons Bypass individual bands so they pass through without processing.

 *Alt-click Solo or Bypass buttons to quickly apply a unique setting to one band.*

Threshold sliders Set the input level at which compression begins. Possible values range from -60 to 0 dB. The best setting depends on audio content and musical style. To compress only extreme peaks and retain more dynamic range, try thresholds around 5 dB below the peak input level; to highly compress audio and greatly reduce dynamic range, try settings around 15 dB below the peak input level.

Input Level meters Measure input amplitude. Double-click the meters to reset peak and clip indicators.

Gain Reduction meters Measure amplitude reduction with red meters that extend from top (minimal reduction) to bottom (maximum reduction).

Gain Boosts or cuts amplitude after compression. Possible values range from -18 to +18 dB, where 0 is unity gain.

Ratio Sets a compression ratio between 1-to-1 and 30-to-1. For example, a setting of 3.0 outputs 1 dB for every 3 dB increase above the compression threshold. Typical settings range from 2.0 to 5.0; higher settings produce the extremely compressed sound often heard in pop music.

Attack Determines how quickly compression is applied when audio exceeds the threshold. Possible values range from 0 to 500 milliseconds. The default, 10 milliseconds, works well for a wide range of audio. Faster settings may work better for audio with fast transients, but such settings sound unnatural for less percussive audio.

Release Determines how quickly compression stops after audio drops below the threshold. Possible values range from 0 to 5000 milliseconds. The default, 100 milliseconds, works well for a wide range of audio. Try faster settings for audio with fast transients, and slower settings for less percussive audio.

Output Gain Boosts or cuts overall output level after compression. Possible values range from -18 to +18 dB, where 0 is unity gain. Double-click the meters to reset peak and clip indicators.

Limiter Applies limiting after Output Gain, at the end of the signal path, optimizing overall levels. Specify Threshold, Attack, and Release settings that are less aggressive than similar band-specific settings. Then specify a Margin setting to determine the absolute ceiling relative to 0 dBFS.



To create extremely compressed audio, enable the Limiter, and then experiment with very high Output Gain settings.

Spectrum On Input Displays the frequency spectrum of the input signal, rather than the output signal, in the multiband graph. To quickly see the amount of compression applied to each band, toggle this option on and off.

Brickwall Limiter Applies immediate, hard limiting at the current Margin setting. (Deselect this option to apply slower soft limiting, which sounds smoother but may exceed the Margin setting.)

Note: *The maximum Attack time for brickwall limiting is 5 ms.*

Link Band Controls Lets you globally adjust the compression settings for all bands, while retaining relative differences between bands.




To temporarily link band controls, hold down Alt+Shift. To reset a control in all bands, hold down Ctrl+Alt+Shift, and click the control.

Normalize effect (Edit View only)

The Amplitude And Compression > 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. If you're sending audio to a mastering engineer, however, normalize audio between -1 and -4 dBFS, providing a cushion for further processing.

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 “Normalize a group of files” on page 258.)

See also

“About process effects” on page 104

“Apply individual effects in Edit View” on page 107


Normalize 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 a percentage.

Normalize L/R Equally Uses both channels of a stereo waveform when the amplification amount is calculated. If this option is deselected, 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.

Tube-modeled Compressor effect

The Amplitude And Compression > Tube-modeled Compressor effect simulates the warmth of vintage hardware compressors. Use this effect to add subtle distortion that pleasantly colors audio.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Tube-modeled Compressor options

Threshold slider Sets the input level at which compression begins. Possible values range from -60 to 0 dB. The best setting depends on audio content and musical style. To compress only extreme peaks and retain more dynamic range, try thresholds around 5 dB below the peak input level; to highly compress audio and greatly reduce dynamic range, try settings around 15 dB below the peak input level.

Input Level meters To the left of the slider, these meters measure input amplitude. Double-click the meters to reset peak and clip indicators.

Gain Reduction meters To the right of the slider, these meters measure amplitude reduction with red bars that extend from top (minimal reduction) to bottom (maximum reduction).

Gain Boosts or cuts amplitude after compression. Possible values range from -18 to +18 dB, where 0 is unity gain.

Ratio Sets a compression ratio between 1-to-1 and 30-to-1. For example, a setting of 3.0 outputs 1 dB for every 3 dB increase above the compression threshold. Typical settings range from 2.0 to 5.0; higher settings produce the extremely compressed sound often heard in pop music.

Attack Determines how quickly compression is applied when audio exceeds the threshold. Possible values range from 0 to 500 milliseconds. The default, 10 milliseconds, works well for a wide range of audio. Faster settings may work better for audio with fast transients, but such settings sound unnatural for less percussive audio.

Release Determines how quickly compression stops after audio drops below the threshold. Possible values range from 0 to 5000 milliseconds. The default, 100 milliseconds, works well for a wide range of audio. Try faster settings for audio with fast transients, and slower settings for less percussive audio.

Delay and echo effects

About delays and echoes

Delays are separate copies of an original signal that reoccur within milliseconds of each other. *Echoes* are sounds that are delayed far enough in time so that you hear each as a distinct copy of the original sound. When reverb or chorus might muddy the mix, both delays and echoes are a great way to add ambience to a track.



To access familiar options from hardware delays, use the Echo effect in Adobe Audition.

See also

“About reverb” on page 152

Analog Delay effect

The Delay And Echo > Analog Delay effect simulates the sonic warmth of vintage hardware delay units. Unique options apply characteristic distortion and adjust the stereo spread. To create discrete echoes, specify delay times of 35 milliseconds or more; to create more subtle effects, specify shorter times.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Analog Delay options

Mode Specifies the type of hardware emulation, determining equalization and distortion characteristics. Tape and Tube reflect the sonic character of vintage delay units, while Analog reflects later electronic delay lines.

Dry Out Determines the level of original, unprocessed audio.

Wet Out Determines the level of delayed, processed audio.

Delay Specifies the delay length in milliseconds.

Feedback Creates repeating echoes by resending delayed audio through the delay line. For example, a setting of 20% sends delayed audio at one-fifth of its original volume, creating echoes that gently fade away. A setting of 200% sends delayed audio at double its original volume, creating echoes that quickly grow in intensity.

Note: When experimenting with extremely high Feedback settings, turn down your system volume.

Trash Increases distortion and boosts low frequencies, adding warmth.

Spread Determines the stereo width of the delayed signal.

Delay effect

The Delay And Echo > Delay effect 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 Chorus or Flanger effects in Adobe Audition, because the delay settings don't change over time.)

By further reducing a delay to between 1 and 14 milliseconds, you can spatially locate a mono sound 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.

See also

"Applying effects in Edit View" on page 106

"Applying effects in Multitrack View" on page 107

"Use effect presets" on page 104

"Add preroll and postroll to effects previews" on page 107

Delay options

Delay Time 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.

Mix Sets the ratio of processed, Wet signal to 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.

Dynamic Delay effect (Edit View only)

The Delay And Echo > 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.

See also

"About process effects" on page 104

"Apply individual effects in Edit View" on page 107

“Control effects settings with graphs” on page 104

“Use effect presets” on page 104

Dynamic Delay 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 Restricts the graph to the length you specify so the effect loops repeatedly over the course of the selection. Selecting this option lets you specify Period (how long the graph should be), Frequency (how fast the loop should repeat), and Total Cycles (how many times you want the graph to loop in the given selection). These settings interact; changing one changes the others. If Loop Graph is deselected, the graph equals the length of the entire selection, so the effect doesn't loop.

Stereo Curve Delay Represents the number of milliseconds one channel is behind the other, creating a stereo phasing sound. Positive values delay the right channel, while negative values delay the left.

Echo effect

The Delay And Echo > Echo 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 sound brighter) to one that is almost totally absorptive (creating echoes that sound darker).

Note: Make sure the audio file is long enough 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 effect.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107


“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Echo options

Left and Right Channel settings Let you define these echo characteristics:

- **Delay Time** Specifies the number of milliseconds, beats, or samples between each echo. For example, a setting of 100 milliseconds results in a 1/10th-second delay between successive echoes.
- **Feedback** 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.
- **Mix** Sets the percentage of echoed (wet) signal to mix with the original (dry) signal in the final output.

 You can create striking stereo echo effects by setting different left and right values for the Delay Time, Feedback, and Mix controls.

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.

Delay Time Units Specifies milliseconds, beats, or samples for the Delay Time setting.

Echo Chamber effect

The Delay And Echo > 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 the audio file is long enough for the echo to end. If the echo is cut off abruptly before it fully decays, undo the Echo Chamber effect, add several seconds of silence by choosing Generate > Silence, and then reapply the effect.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104


“Add preroll and postroll to effects previews” on page 107

Echo Chamber 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** Specifies the distance 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 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.*

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.

Sum 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.

 *To give mono audio ambiance, convert it to stereo (choose Edit > Convert Sample Type), and then apply the Echo Chamber effect. Set the left microphone one or two feet from the right microphone to simulate a listener's ears, an especially effective technique for headphones.*

Multitap Delay effect

The Delay And Echo > Multitap Delay effect 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. A single delay unit is much the same as the Echo function, but with a slightly different filtering setup. (A delay unit uses two sliding bands with variable cutoff points instead of eight bands of filtering.)

Experiment to achieve very interesting effects. If one delay unit is placed inside another, then the echo occurs more than once. As audio travels down the delay line, it can be fed back into the line with varying offset, feedback, and filter settings.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Multitap Delay 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:

- **Delay 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 as if they were at 100 and 400. The difference between the offsets is what is important.

- **Delay Time** 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.

- **Gain** Sets the amount of filtering. Gain 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.

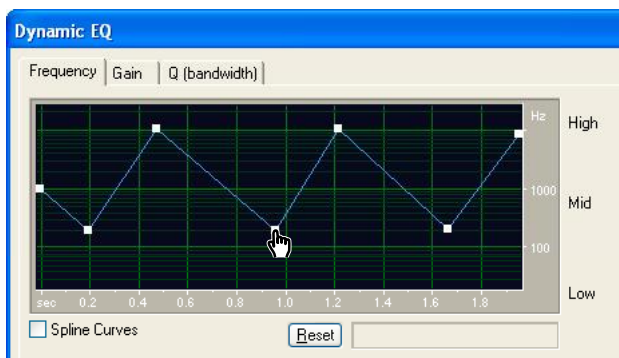
Channels menu Provides the following options for use with stereo source audio. To spread out the effect in the stereo domain, at least one delay unit in the group should be cross-channel or single-channel:

- **Left Only** Adds delay to only the left channel.
- **Right Only** Adds delay to only the right channel.
- **Discrete Stereo** Adds delay to both channels.
- **Stereo Swap** Maps the delay for the left channel to the right channel and vice-versa.
- **From Right To Left** Delays audio from the right channel and adds it to the left channel.
- **From Left To Right** Delays audio from the left channel and adds it to the right channel.
- **From Center To Surround** Delays audio common to both channels and inverts phase of the left-channel delay relative to the right-channel delay.

Filter and equalizer effects

Dynamic EQ effect (Edit View only)

The Filter And EQ > 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 Dynamic EQ effect (Rhythmic Sweep preset)

See also

“About process effects” on page 104

“Apply individual effects in Edit View” on page 107

“Control effects settings with graphs” on page 104

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Dynamic EQ 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 frequency bandwidth affected by the Dynamic EQ effect. The horizontal ruler (*x*-axis) represents time, and the vertical ruler (*y*-axis) represents bandwidth, measured in Q values. (As Q values increase, bandwidth narrows.)

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 Restricts the graph to the length you specify so the effect loops repeatedly over the course of the selection. Selecting this option enables a menu for specifying Period (how long the graph should be), Frequency (how fast the loop should repeat), Tempo (how many beats per minute the loop should have), and Total Cycles (how many times you want the graph to loop in the given selection). These settings interact; changing one changes the others. If Loop Graph is deselected, the graph equals the length of the entire selection, so the effect doesn't loop.

FFT Filter effect (Edit View only)

The graphic nature of the Filter And EQ > FFT Filter effect makes it easy to draw curves or notches that reject or boost specific frequencies. (For a definition, see “FFT” on page 275.)

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.

 For optimal results, filter 32-bit samples so you can perform multiple transforms without losing audible data. If the source audio is 8- or 16-bit, convert it to 32-bit, filter it, and convert back to 8- or 16-bit with dithering.



FFT Filter graph (De-Esser preset)

See also

“About process effects” on page 104

“Apply individual effects in Edit View” on page 107

“Control effects settings with graphs” on page 104

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

FFT Filter 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.

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 the transition percentage along the *y*-axis (where 0% represents the initial filter, and 100% represents the final filter).

The Graph Response At Point option causes the lower readout to change in response to cursor position in the graph above, visually representing the filter at any point in the transition. Depending on the position you select, you can specify a morphing transition or a linear transition.

Graphic Equalizer effect

The Filter And EQ > 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.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Graphic Equalizer options


Bands buttons 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 section of the Graphic Equalizer window are identical.

Band Determines the band to be modified with the Gain value.

Gain Sets the exact boost or attenuation (measured in decibels) for the chosen band.

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.)


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.*

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.

Notch Filter effect

The Filter And EQ > 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 remove excessive sibilance (shrill “ess” sounds), use the De-Essing preset.*

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Notch Filter 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.

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.

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.

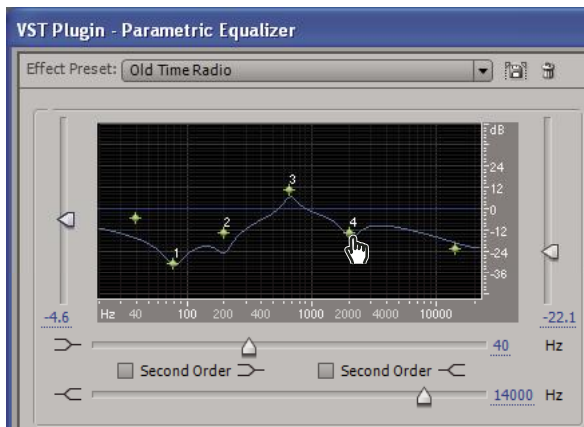
 *Use no more than 30 dB attenuations for a Narrow setting, no more than 60 dB for Very Narrow, and no more than 90 dB for Super Narrow. Greater attenuation can remove a wide range of neighboring frequencies.*

Ultra-Quiet Virtually eliminates noise and artifacts, but requires significant processing. This option is audible only on state-of-the-art, professional monitoring systems.

Parametric Equalizer effect

The Filter And EQ > Parametric Equalizer effect 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)

See also

“Applying effects in Edit View” on page 106

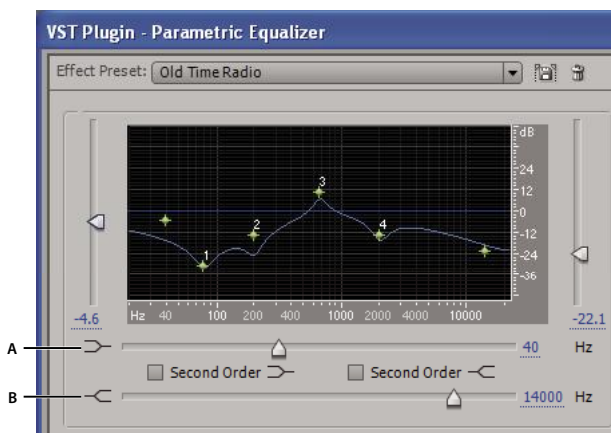
“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Parametric Equalizer options

Graph Shows 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).



Parametric Equalizer effect with Old Time Radio preset:
 A. Low Shelf Cutoff and Level B. High Shelf Cutoff and Level

Low Shelf Cutoff and Level Determines affected low frequencies and their amplitude. Set the cutoff frequency with the slider directly below the graph, and increase or decrease bass with the slider at the left of the graph. Use the Low Shelf Cutoff to reduce low-end rumble, hum, or other unwanted low-frequency sounds.

High Shelf Cutoff and Level Determines affected high frequencies and their amplitude. Set the cutoff frequency with the second slider below the graph, and increase or decrease treble with the slider at the right of the graph. Use the High Shelf Cutoff to reduce hiss, amplifier noise, and the like.

Second Order buttons Adjust the low and high shelves by 12 dB per octave. (The default, first-order filters adjust the shelves by 6 dB per octave.)

Frequency Bands 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 in the text box below each slider.

Amplitude sliders Set the boost or attenuation for each frequency band.

Ultra-Quiet Virtually eliminates noise and artifacts, but requires significant processing. This option is audible only on state-of-the-art, professional monitoring systems.

Constant Q, Constant Width 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.

Q / Width Controls the width of the affected frequency band, measured in either Q or Width values. Low Q values (or high Width values) affect a larger range of frequencies. Very high Q values (close to 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.

Gain Compensates for an overall volume level that might be too loud or too soft after you adjust the EQ settings.

Quick Filter effect (Edit View only)

The Filter And EQ > Quick Filter effect 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.

See also

“About process effects” on page 104

“Apply individual effects in Edit View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Quick Filter options

Lock To These Settings Only Determines whether settings are fixed or varied.

If this option is selected, the entire selected range is equalized at the settings shown.


If this option is deselected, 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 is deselected. Click this tab to specify the initial EQ settings.


Final Settings tab Appears if Lock To These Settings Only is deselected. 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.*

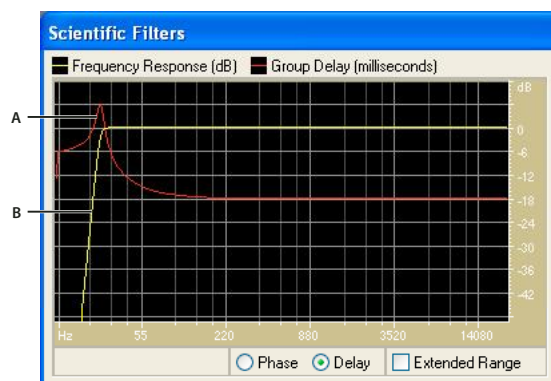
Lock L/R Lets you adjust channels together or separately. Select this option to adjust the channels together, maintaining the same settings for each. Leave this option unselected to adjust each channel separately.

 *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.*

Scientific Filters effect (Edit View only)

The Filter And EQ > 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: Bessel, Butterworth, 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)

See also

“About process effects” on page 104

“Apply individual effects in Edit View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Scientific Filters 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 the 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.

Modulation effects

Chorus effect

The Modulation > Chorus effect simulates several voices or instruments played at once by adding multiple short delays with a small amount of feedback. The result is lush, rich sound. You can use Chorus to enhance a vocal track or add stereo spaciousness to mono audio. 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, making each voice (or layer) sound distinct from the original by slightly varying timing, intonation, and vibrato. The Feedback setting lets you add extra detail to the result.



To achieve the best results with mono files, convert them to stereo before applying the Chorus effect.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Chorus options

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.

- **Voices** Determines the number of simulated voices.

Note: As you add more voices, the sound becomes richer and richer—but processing time also increases.

- **Delay Time** 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, like a tape being eaten by a cassette deck.
- **Delay Rate** Determines the time the delay takes to cycle from zero to the 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 separate voices). For example, a value of 2 Hz would vary the delay from zero to the maximum and back twice per second (simulating a pitch vibrato at twice per second). 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** Adds a percentage of processed voices back into the effect input. Feedback can give a waveform an extra echo or reverb effect. A little feedback (less than 10%) can provide extra richness, depending on the delay and vibrato settings. Higher settings produce more traditional feedback, a loud ringing which can get loud enough to clip the signal. Sometimes this clipping is a desired effect, as in the Flying Saucers preset, which generates the 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 voice 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.
- **Modulation 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 depth may be unnoticeable unless the Modulation 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.
- **Modulation Rate** Determines the maximum rate at which amplitude changes occur. 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).

- **Highest Quality** Ensures the best quality results. Increasing the quality, however, increases the processing time for previewing and applying the effect.

Stereo Width Determines where the individual voices are placed in the stereo field and how the original stereo signal is interpreted. These options are active only when you work with stereo files:

- **Average Left & Right Channel Input** Averages the original left and right channels. If the option is deselected, the channels are kept separate to preserve the stereo image. (For example, spatial binaural cues such as those that exist in reverberated audio or live stereo recordings are preserved.) Leave this option deselected if the stereo source audio was originally monophonic—it won't have any effect other than increasing 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 when you listen through headphones. Deselect this option for audio that will be played through speakers. (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 you don't add cues, however, no output is sent to the right channel, so greater separation is heard when listening through speakers.)
- **Stereo 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 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 Level Sets the ratio of original (Dry) signal to chorused (Wet) signal. Extremely high settings may cause clipping.



In Multitrack View, you can vary the Wet level over time with automation lanes. (See “Automating track settings” on page 204.) This technique is handy for emphasizing vocal or instrumental solos. Keep the Dry level near 100% and reduce the Wet level to around 30% to give a solo a “backup chorus.”

Flanger effect

Flanging is an audio effect caused by mixing a varying, short delay in roughly equal proportion to the original signal. It was originally achieved by sending an identical audio signal to two reel-to-reel tape recorders, and then pressing the flange of one reel to slow it down. Combining the two resulting recordings produced a phase-shifted, time-delay effect, characteristic of psychedelic music of the 1960s and 1970s. The Modulation > Flanger effect lets you create a similar result by slightly delaying and phasing a signal at specific or random intervals.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Flanger options

Initial Delay Time 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 Delay Time 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° 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.

Modulation 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.

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 Effects** 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.

Mix Adjusts the mix 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.

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 Modulation > Sweeping Phaser effect sweeps a notch- or boost-type filter back and forth across a center frequency.

Like flangers, phasers add a slight delay, but they also phase-shift frequencies over time. On stereo files, the stereo image can be dramatically altered to create some remarkably interesting sounds.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Sweeping Phaser 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.

Modulation 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.

Modulation 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.

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.

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.


Output 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.


Restoration effects

Techniques for restoring audio

You can fix a wide array of audio problems by combining two powerful features. First, use Spectral Frequency Display to visually identify and select ranges of noise or individual artifacts. (See “Select spectral ranges” on page 73 and “Select artifacts and repair them automatically” on page 74.) Then, use restoration effects to fix problems like the following:

- Crackle from wireless microphones or old vinyl records. (See “Automatic Click Remover effect” on page 142.)
- Background noise like wind rumble, tape hiss, or 60-cycle hum. (See “Adaptive Noise Reduction effect” on page 141.)
- Phase misalignment from poorly placed stereo microphones or aging analog tapes. (See “Automatic Phase Correction effect” on page 142.)

 *The VST-based restoration effects above, which are available in both Edit View and Multitrack View, quickly address common audio problems. For unusually noisy audio, however, consider using offline effects unique to Edit View, such as Click/Pop Eliminator, Hiss Reduction, and Noise Reduction.*



Selecting various types of noise in Spectral Frequency Display
A. Hiss B. Crackle C. Rumble

See also

“About Spectral Frequency Display” on page 66

“Edit audio clips from After Effects” on page 227

Adaptive Noise Reduction effect

The Restoration > Adaptive Noise Reduction effect quickly removes variable broadband noise such as background sounds, rumble, and wind. Because this effect is VST-based, you can combine it with other effects in the Mastering Rack and Effects Rack and apply it to tracks in Multitrack View. By contrast, the standard Noise Reduction effect is available only as an offline process in Edit View. That effect, however, is sometimes more effective at removing constant noise, such as hiss or hum.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Adaptive Noise Reduction options

Reduce Noise By Determines the level of noise reduction. Values between 6 and 30 dB work well. To reduce bubbly background effects, enter lower values.

Noisiness Indicates the percentage of original audio that contains noise.

Fine Tune Noise Floor Adjusts the noise floor above or below the automatically calculated floor.

Signal Threshold Adjusts the threshold of desirable audio above or below the automatically calculated threshold.

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.

Broadband Preservation Retains desirable audio in specified frequency bands between found artifacts. A setting of 100 Hz, for example, ensures that no audio is removed 100 Hz above or below found artifacts.

FFT Size Determines how many individual frequency bands are analyzed. Choose a high setting to increase frequency resolution; choose a low setting to increase time resolution. High settings work well for artifacts of long duration (like squeaks or 60-Hz hum), while low settings better address transient artifacts (like clicks and pops).

High Quality Mode Performs slower processing but achieves superior results.

Automatic Click Remover effect

To quickly remove crackle and static from vinyl recordings, use the Restoration > Automatic Click Remover effect. You can 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. And because the Automatic Click Remover is VST-based, you can combine it with other effects in the Mastering Rack and Effects Rack and apply it to tracks in Multitrack View.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Automatic Click Remover options

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 30.

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 16.

Automatic Phase Correction effect

The Restoration > Automatic Phase Correction effect addresses azimuth errors from misaligned tape heads, stereo smearing from incorrect microphone placement, and many other phase-related problems.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104


“Add preroll and postroll to effects previews” on page 107

Automatic Phase Correction options

Global Time Shift Activates the Left and Right Channel Shift sliders, which let you apply a uniform phase shift to all selected audio.

Auto Align Channels and Auto Center Panning Align phase and panning for a series of discrete time intervals, which you specify using the following options:


- **Time Resolution** Specifies the number of milliseconds in each processed interval. Smaller values increase accuracy; larger ones increase performance.
- **Responsiveness** Determines overall processing speed. Slow settings increase accuracy; fast settings increase performance.
- **Channel** Specifies the processed channels.
- **Analysis Size** Specifies the number of samples in each analyzed unit of audio.


 *For the most precise, effective phase correction, use the Auto Align Channels option. Enable the Global Time Shift sliders only if you are confident that a uniform adjustment is necessary, or if you want to manually animate phase correction in Multitrack View.*

Click/Pop Eliminator effect (Edit View only)

The Restoration > Click/Pop Eliminator effect detects and removes clicks and pops. Like the Automatic Click Remover, 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 customize settings to repair other transient artifacts such as vocal plosives or radio static.

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).

 *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 Preferences dialog box.) Most clicks appear as bright vertical bars that extend from the top to the bottom of the waveform display.*



Click/Pop detection graph

A. Level of detected clicks and pops B. Level of rejected clicks and pops

See also

“About process effects” on page 104

“Apply individual effects in Edit View” on page 107

“Use effect presets” on page 104

“Control effects settings with graphs” on page 104

Click/Pop Eliminator 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.



To quickly remove clicks and pops, select a preset, and click Auto Find All Levels. Then fine-tune settings if necessary.

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 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 "Mix audio data when pasting" on page 77.

Clip Restoration effect (Edit View only)

The Restoration > 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.

See also

“About process effects” on page 104

“Apply individual effects in Edit View” on page 107

“Use effect presets” on page 104

Clip Restoration 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 “Change the bit depth of a file” on page 99.) 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.*

Hiss Reduction effect (Edit View only)

The Restoration > 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.

 To reduce other types of noise that have a wide frequency range, try the Noise Reduction effect. (See “Noise Reduction effect (Edit View only)” on page 149.)



Using the Hiss Reduction graph to adjust the noise floor

See also

“About process effects” on page 104

“Apply individual effects in Edit View” on page 107

“Use effect presets” on page 104

“Control effects settings with graphs” on page 104

“Add preroll and postroll to effects previews” on page 107

Hiss Reduction 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 high, some hiss may remain after processing.

If the value is too low, other background artifacts might be heard.

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 high (above 90%), unnaturally long tails and reverbs might be heard.

If the value is too low, background bubbly effects might be heard, and music might sound artificial.

Reduce Hiss By Sets the level of hiss reduction for audio below the noise floor.

With higher values (especially above 20 dB) dramatic hiss reduction can be achieved, but the remaining audio might become distorted.

With lower values, not as much noise is removed, and the original audio signal stays relatively undisturbed.


Remove Hiss, Keep Only Hiss Removes hiss or removes all audio except for hiss.

Noise Reduction effect (Edit View only)


The Restoration > 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.


See also

“About process effects” on page 104

“Control effects settings with graphs” on page 104

Apply the Noise Reduction effect

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 “Select spectral ranges” on page 73.)

2 Choose Effects > Restoration > Capture Noise Reduction Profile.

3 In the Main panel, select the range from which you want to remove noise.

4 Choose Effects > Restoration > Noise Reduction.


5 Set the desired options.

Noise Reduction 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 any noise profile previously saved from Adobe Audition in FFT format. However, you can apply noise profiles only to identical sample types. (For example, you can't apply a 22 kHz, mono, 16-bit profile to 44kHz, stereo, 8-bit samples.)

Note: Because noise profiles are so specific, a profile for one type of noise won't produce good results with other types. If you regularly remove similar noise, however, a saved profile can greatly increase efficiency.

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 Reset), 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 results, 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** Determines the level of noise reduction. Values between 6 and 30 dB work well. To reduce bubbly background effects, enter lower values.
- **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.

Reverb effects

About reverb

In a room, sound bounces off the walls, ceiling, and floor on the way to your ears. All these reflected sounds reach your ears so closely together that you don't perceive them as separate echoes, but as a sonic ambience that creates an impression of space. This reflected sound is called *reverberation*, or *reverb* for short. With Adobe Audition, you can use reverb effects to simulate a variety of room environments.

 For the most flexible, efficient use of reverb in Multitrack View, add reverb effects to buses, and set reverb output levels to 100% Wet. Then, route tracks to these buses, and use sends to control the ratio of dry to reverberant sound.

See also

“About delays and echoes” on page 122

“Routing audio to buses, sends, and the Master track” on page 186

Convolution Reverb effect

The Reverb > Convolution Reverb effect reproduces rooms ranging from coat closets to concert halls. Convolution-based reverbs use *impulse* files to simulate acoustic spaces. The results are incredibly realistic and life-like.

Because Convolution Reverb requires significant processing, you may hear clicks or pops when previewing it on slower systems. These artifacts disappear after you apply the effect.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Convolution Reverb options

Impulse Specifies a file that simulates an acoustic space. Click Load to add a custom impulse file.

Mix Controls the ratio of original to reverberant sound.

Room Size Specifies a percentage of the full room defined by the impulse file. The larger the percentage, the longer the reverb.

Damping LF Reduces low-frequency, bass-heavy components in reverb, avoiding muddiness and producing a clearer, more articulate sound.

Damping HF Reduces high-frequency, transient components in reverb, avoiding harshness and producing a warmer, lush sound.

Pre-Delay Determines how many milliseconds the reverb takes to build to maximum amplitude. To produce the most natural sound, specify a short pre-delay of 0–10 milliseconds. To produce interesting special effects, specify a long pre-delay of 50 milliseconds or more.

Width Controls the stereo spread. A setting of 0 produces a mono reverb signal.

Gain Boosts or attenuates amplitude after processing.

Full Reverb effect

The Reverb > Full Reverb effect is convolution-based, avoiding ringing, metallic, and other artificial sounding artifacts.

This effect offers some unique options, such as Perception, which simulates room irregularities, Left/Right Location, which places the source off-center, and Room Size and Dimension, which help you realistically simulate rooms that you can customize. To simulate wall surfaces and resonance, you can change the reverb's frequency absorption by using a three-band, parametric EQ in the Coloration section.

Note: *The Reverb effect demands significant processing; for real-time, multitrack use, either freeze this effect or replace it with Studio Reverb. (See “Freeze effects and improve performance” on page 109.)*

When you change reverb settings, this effect creates a new *impulse* file, which simulates the acoustic environment you specify. This file can be several megabytes in size, requiring a few seconds to process, so you might have to wait before hearing a preview. The results, however, are incredibly realistic and easy to tailor.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

“Control effects settings with graphs” on page 104


Reverb Settings options

Click Reverb Settings to display the following options:


Decay Time 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 to allow for decaying into the background noise level.

Pre-Delay 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 pre-delay times of 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. Lower values produce more distinct echoes because the initial echo density is lighter, but the density builds over the life of the reverb tail.

 *Bouncy echo 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.*

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 Output Level section adjusts the original signal to sound as if it were coming from the same location. Very nice effects are possible with singers being very slightly off center, 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 Room Size Sets an appropriate decay length and pre-delay time to match the specified room size and produce a more convincing reverb. If desired, you can then fine-tune the decay length and pre-delay.

Coloration options


Click Coloration to display the following options for filtering the reverb:

Amplitude sliders Provide a quick way to adjust reverb volume. They are located to the right of the graph and specify the low shelf, mid band, and high shelf amplitude, 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 sliders Located below the graph, these specify the corner 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 the width of the mid band. Higher values affect a narrower range of frequencies, and lower 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.*

Decay 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.

Output Level options

The Output Level section provides the following options for balancing wet and dry signals:


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.

Wet (reverb) Controls the volume of the dense layer of reverberant. Like the Wet setting for early reflections, this option should be fine-tuned to give a pleasing mix. If it's too loud, the reverb sounds very unnatural. The balance between reverb and the original signal gives the impression of distance, so increase the reverb volume to make the source sound farther away.

Wet (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.

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.

Sum Left And Right Inputs 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 mono sample).*

Reverb effect

The Reverb > Reverb effect simulates acoustic spaces with convolution-based processing. It produces both early reflections and irregular echoes that are so closely spaced that they're perceived as a single decaying sound.

The Reverb effect can create a wide range of high-quality 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.

Relative to the Reverb effect, the Full Reverb effect provides more options and better audio rendering. For quick adjustments, however, you may prefer the reduced options set of the Reverb effect.

Note: The Reverb effect demands significant processing; for real-time, multitrack use, either freeze this effect or replace it with Studio Reverb. (See "Freeze effects and improve performance" on page 109.)

See also

"Applying effects in Edit View" on page 106


"Applying effects in Multitrack View" on page 107

"Use effect presets" on page 104

"Add preroll and postroll to effects previews" on page 107

Reverb options


Decay Time 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.

 *To simulate rooms that have both echoes and reverb, use the Echo effect first to establish the size of the room, and then use the Reverb effect to make the sound more natural. A Decay Time as little as 300 milliseconds can add perceived spaciousness to dry sound.*


Pre-Delay 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.

Diffusion 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.*

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.*

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).

Sum Left And Right Inputs 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.

Studio Reverb effect

Like the other reverb effects, the Reverb > Studio Reverb effect simulates acoustic spaces. It is faster and less processor-intensive than the other reverb effects, however, because it isn't convolution-based. As a result, you can make real-time changes quickly and effectively in Multitrack View, without freezing effects on a track.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Studio Reverb options

Room Size Sets the room size.

Decay Adjusts the amount of reverberation decay in milliseconds.

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.

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.

Dry Sets the percentage of source audio to output with the effect.

Wet Sets the percentage of reverb to output.

Special effects


Convolution effect (Edit View only)

The Special > Convolution effect multiplies selected audio samples by samples in an *impulse* (.imp) file. Impulses function like amplitude maps. For example, if you apply an impulse of a single full-volume sample, the selected audio data will be unchanged. Should the impulse be at half volume, however, the selected audio data will be reduced to half volume. If several such impulses occur over time, each with descending amplitude, the selected audio data will gradually and rhythmically become lower in volume.

This effect essentially applies the sound of one audio file to another. The result can resemble filtering, echoing, phase shifting, or any combination of these effects. 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 create a stereo impulse in your favorite cathedral, and you convolute it with any mono audio, the result sounds as if that audio were played in that cathedral. (To create such an impulse, you would place a stereo microphone in the cathedral, generate a loud impulsive noise like a hand clap, and record the noise and reverberation. If you apply the resulting impulse to other audio, it sounds as though it was generated at the location of the hand clap relative to the microphone.)

To get a feel for the Convolution effect, open and play with some of the sample impulse files in the Imps folder in the Adobe Audition program folder.

 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.

See also

“About process effects” on page 104

“Apply individual effects in Edit View” on page 107

Convolution 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.

Distortion effect

Use the Special > Distortion 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.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Control effects settings with graphs” on page 104

Distortion options

Symmetric Changes the buttons in the dialog box from (Symmetric)/(Symmetric) to Positive/Negative.

(Symmetric)/(Symmetric) buttons Create identical positive and negative curves. You cannot choose one over the other.

Positive/Negative buttons Let you specify separate distortion curves for positive and negative sample values. Select the tab corresponding to the distortion curve you want to display.

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.

Guitar Suite effect

The Special > Guitar Suite effect applies a series of processors that optimize and alter the sound of guitar tracks. The Compressor stage reduces dynamic range, producing a tighter sound with greater impact. Filter, Distortion, and Box Modeler stages simulate common effects that guitarists use to create expressive, artistic performances.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Guitar Suite options

Compressor Reduces dynamic range to maintain consistent amplitude and help guitar tracks stand out in a mix.

Filter Simulates guitar filters ranging from resonators to talk boxes. Choose an option from this menu, and then set options below:

- **Type** Determines which frequencies are filtered. Specify Lowpass to filter high frequencies, Highpass to filter low frequencies, or Bandpass to filter frequencies above and below a center frequency.
- **Freq** Determines the cutoff frequency for Lowpass and Highpass filtering, or the center frequency for Bandpass filtering.
- **Resonance** Feeds back frequencies near the cutoff frequency, adding crispness with low settings and whistling harmonics with high settings.

Distortion Adds a sonic edge often heard in guitar solos. To change the distortion character, choose an option from the Type menu.


Box Modeler Simulates various amplifiers that guitarists use to create unique tones.

Mix Controls the ratio of original to processed audio.

Mastering effect

Mastering describes the complete process of optimizing audio files for a particular medium, such as radio, video, CD, or the web. In Adobe Audition, you can quickly master audio with the Special > Mastering effect.

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 computer speakers that poorly reproduce bass sounds. To compensate, you can boost bass frequencies during the equalization stage of the mastering process.

 *Adobe Audition provides a default workspace that's optimized for mastering. Choose Window > Workspace > Mastering And Analysis.*

See also

“Techniques for restoring audio” on page 141

“Normalizing groups of files” on page 258

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107


“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Mastering options

Equalizer Adjusts the overall tonal balance.

- **Graph** Shows 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).

 *Drag control points in the graph to visually adjust the settings below.*

- **Low Shelf and High Shelf Enable** Activate shelving filters at either end of the frequency spectrum.
- **Peaking Enable** Activates a peaking filter in the center of the frequency spectrum.
- **Hz** Indicates the center frequency of each frequency band.
- **dB** Indicates the level of each frequency band.
- **Q** Controls the width of the affected frequency band. Low Q values (up to 3) affect a larger range of frequencies and are best for overall audio enhancement. High Q values (6–12) affect a very narrow band and are ideal for removing a particular, problematic frequency, like 60-Hz hum.

Reverb Adds ambience. Drag the Amount slider to change the ratio of original to reverberant sound.

Exciter Exaggerates high-frequency harmonics, adding crispness and clarity. Mode options include Retro for light distortion, Tape for bright tone, and Tube for quick, dynamic response. Drag the Amount slider to adjust the level of processing.

Widener Adjusts the stereo image. Drag the Width slider to the left to narrow the image and increase central focus. Drag the slider to the right to expand the image and enhance spatial placement of individual sounds.

Loudness Maximizer Applies a limiter that reduces dynamic range, boosting perceived levels. A setting of 0% reflects original levels; 100% applies maximum limiting.

Output gain Determines output levels after processing. For example, to compensate for EQ adjustments that reduce overall level, boost the output gain.

Stereo imagery effects

Changing stereo imagery

Adobe Audition provides several effects that let you change the apparent location, or *stereo imagery*, of sounds coming from the speakers. For example, 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: Stereo imagery effects work only on stereo files.

Binaural Auto-Panner effect (Edit View only)

The Stereo Imagery > Binaural Auto-Panner effect rotates sound spatially from left to right in a seemingly circular pattern. This effect delays either the left or right channel so sounds reach each ear at different times, tricking the brain into thinking sounds are coming from either side.

See also

“About process effects” on page 104

“Apply individual effects in Edit View” on page 107

“Control effects settings with graphs” on page 104

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Binaural Auto-Panner options

Frequency graph Represents time along the *x*-axis (horizontal ruler) and frequency along the *y*-axis (vertical ruler). 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.

Bottom Graph Settings, Top Graph Settings Control the low and high binaural frequencies, represented by the bottom and top of the graph:

- **Pan Cycling Rate** Determines 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 moves it in 10 seconds.
- **Intensity** Controls the degree of 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).

Center Channel Extractor effect

The Stereo Imagery > 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 vocals, bass, or kick drum, or you can remove any of them to create a karaoke mix.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Center Channel Extractor options

Extract Audio From Specifies the phase degree, pan percentage, and delay time for the audio you want to extract or remove. Set this option to Center (0°) 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°) 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, Pan, and Delay. The former two settings can range from -100% (hard left) to 100% (hard right). The latter setting ranges from -5 to +5 milliseconds.

Frequency Range Sets the range you want to extract or remove. Predefined ranges include Male Voice, Female Voice, Bass, and Full Spectrum. 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.

Amplitude Discrimination and Amplitude Bandwidth Sums the left and right channels, and creates a 180°-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 Bandwidth 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 Bandwidth settings between 1 and 20 work 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.

FFT Size Specifies the Fast Fourier Transform size, 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 FFT windows 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) between each FFT window. Values between 10 and 50 milliseconds usually work best, but higher overlay settings may require a different value.

Window Width Specifies the width (measured as a percentage) of each FFT window. Values of 30% to 100% work well.

Channel Mixer effect

The Stereo Imagery > 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.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Channel Mixer 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 results in a new left channel that contains equal audio from both the current L and R channels. In contrast, an L value of 0 and a R value of 100 results in a new left channel that contains only audio from the current right channel.

New Right Channel Determines the percentage of the current left and right channels to mix into the new right channel. For examples, see above.

Invert Inverts a channel's phase polarity (that is, turns peaks into valleys and valleys into 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.

Doppler Shifter effect (Edit View only)

The Stereo Imagery > 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.

See also

“About process effects” on page 104

“Apply individual effects in Edit View” on page 107

“Use effect presets” on page 104

Doppler Shifter options

Note: Unlike many 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.

Path Type Defines which path the sound source appears to take. Depending on the path type, a different set of options is available.

Straight Line 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 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.

Graphic Panner effect

The Stereo Imagery > Graphic Panner effect combines panning, expansion, and phase adjustment, giving you the power to dramatically alter the stereo image. A unique graph for each adjustment type lets you visually evaluate and change the stereo position of each frequency. To automatically create graph curves that compensate for off-center audio or reveal current stereo position, select Auto Center or Learn. With either option, the graph is updated dynamically as you play audio.

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Graphic Panner options

Pan, Expand, and Phase Curve buttons Display a unique graph and options for each curve type.

Graph Visually indicates the adjustment at each frequency. Click and drag to create and move control points.

Frequency Scale Determines how frequencies are arranged in the horizontal ruler. A Linear scale provides equal weight to each frequency; a Logarithmic scale provides more low-frequency detail, better reflecting the emphasis of human hearing.

Auto Center Pan or Phase Analyzes selected or played audio, and creates a graph curve that compensates for current panning or phase position, placing each frequency in the center of the stereo field.

Learn Pan, Expand, or Phase Analyzes selected or played audio, and creates a graph curve that reflects current panning, expansion, or phase position for each frequency.



The Learn options provide a helpful starting point for your own adjustments.

Panning Method Pans either by raising and lowering amplitude in each channel, or by moving audio across the stereo field.

Auto-Cycle Modulates panning, expansion, and phase adjustments at the frequency or beats you specify, moving audio back and forth across the stereo field.

FFT Size Determines the size of each piece of processed audio. Choose a high setting to adjust frequencies more accurately but time location less accurately. Or, choose a low setting to do the opposite.

Pan/Expand effect (Edit View only)

The Stereo Imagery > 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.

See also

“About process effects” on page 104

“Apply individual effects in Edit View” on page 107

“Control effects settings with graphs” on page 104

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Pan/Expand options

Center Channel Pan graph 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 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).

Stereo Expander effect

Like the Pan/Expand effect, the Stereo Imagery > Stereo Expander effect positions and expands the stereo image. Because the Stereo Expander is VST-based, however, you can combine it with other effects in the Mastering Rack and Effects Rack. In Multitrack View, you can also vary the effect over time by using automation lanes. (See “Automating track settings” on page 204.)

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Stereo Expander options

Center Channel Pan Positions the center of the stereo image anywhere from hard left (-100%) to hard right (100%).

Stereo Expand Expands the stereo image from Narrow/Normal (0) to Wide (300). Narrow/Normal reflects the original, unprocessed audio.

Stereo Field Rotate VST effect

Like the similar process effect, the Stereo Imagery > Stereo Field Rotate VST effect rotates the left and right channels. However, you can combine the VST version with other effects in the Mastering Rack and Effects Rack. In Multitrack View, you can also vary the effect over time by using automation lanes. (See “Automating track settings” on page 204.)

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Stereo Field Rotate VST options

Rotation Sets the degree of rotation.

Invert Left/Right Reverses the rotation.

Range menu Sets the range of rotation from 45° to 360°.

When specifying Rotation and Range settings, keep the following in mind:

- If you rotate the stereo field 180°, both the left and right channels are 180° out of phase, inverting the entire waveform.
- At 90° right, only the right channel is inverted; at 90° left, only the left channel is inverted.
- A rotation of 45° produces results identical to the Pan/Expand effect, panning centered audio to the left or right, while panning surrounding audio in the opposite direction.

Stereo Field Rotate process effect (Edit View only)

The Stereo Imagery > Stereo Field Rotate process rotates the left and right channels, changing where instruments and other sources appear to be located. By manipulating the Rotation graph, you can rotate sources over time.

See also

“About process effects” on page 104

“Apply individual effects in Edit View” on page 107


“Control effects settings with graphs” on page 104

“Use effect presets” on page 104

Stereo Field Rotate process options

Rotation graph 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) at any point in time. The *x*-axis (horizontal ruler) of the graph shows the waveform’s timeline, while the *y*-axis (vertical ruler) displays the number of degrees off center for the left and right channels.

Invert Left/Right Reverses the graph so that an upward line spins the stereo field clockwise instead of counter-clockwise.

 *With the Phase Analysis window (Window > Phase Analysis), you can watch the stereo field rotate. For example, if you create a graph that rotates audio to the right, the phase rotates clockwise. But if you also select Invert Left/Right, the phase rotates counter-clockwise.*

Range Sets the range of the *y*-axis from 45° to 360°.

Note: *If you rotate the stereo field 180°, both the left and right channels are 180° out of phase, inverting the entire waveform. At 90° right, only the right channel is inverted. At 90° left, only the left channel is inverted. A rotation of 45° produces results identical to the Pan/Expand effect, panning centered audio to the left or right, while panning surrounding audio in the opposite direction.*

Loop Graph Restricts the graph to the length you specify so the effect loops repeatedly over the course of the selection. Selecting this option enables a menu for specifying Period (how long the graph should be), Frequency (how fast the loop should repeat), Tempo (how many beats per minute the loop should have), and Total Cycles (how many times you want the graph to loop in the given selection). These settings interact; changing one changes the others. If Loop Graph is deselected, the graph equals the length of the entire selection, so the effect doesn’t loop.

Time and pitch manipulation effects

Pitch Bender effect (Edit View only)

The Time And Pitch > Pitch Bender effect varies the pitch of audio over time. Use the graph to draw a tempo and create smooth tempo changes or other effects, like a vinyl record speeding up or slowing down.

See also

“About process effects” on page 104

“Apply individual effects in Edit View” on page 107

“Control effects settings with graphs” on page 104

“Use effect presets” on page 104

Pitch Bender 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.

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.

Pitch Correction effect (Edit View only)

The Time And Pitch > Pitch Correction effect provides two ways to adjust the pitch adjustments of vocals or solo instruments. Automatic mode analyzes the audio content and automatically corrects the pitch based on the key you define, without making you 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 is most effective with audio that contains a periodic signal (that is, audio with one note at a time, such as saxophone, violin, or vocal parts). Nonperiodic audio, or any audio with a high noise floor, can disrupt the effect's ability to detect the incoming pitch, resulting in incomplete pitch correction.

See also

“About process effects” on page 104

“Apply individual effects in Edit View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Automatic options

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.

Manual options

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 has two main components:

- **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).

- **Readout below the graph** Shows the position of a control point as you drag it.

Vertical scroll 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 scroll bar Lets you scroll to a different time range of the waveform in both graphs. Zoom in or out by dragging the edge of the slider.



You can also zoom by right-clicking the sliders and choosing Zoom Out or Zoom Out Full.

Zoom buttons Zoom in and out of the Pitch Edit graph. 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 isn't affected by vertical zooming, so it always retains the same vertical size.)

Pitch Shifter effect

The Time And Pitch > Pitch Shifter effect changes musical pitch. Unlike the other pitch manipulation effects, however, Pitch Shifter is VST-based, so you can combine it with other effects in the Mastering Rack or Effects Rack. In Multitrack View, you can also vary pitch over time by using automation lanes. (See “Automating track settings” on page 204.)

See also

“Applying effects in Edit View” on page 106

“Applying effects in Multitrack View” on page 107

“Use effect presets” on page 104

“Add preroll and postroll to effects previews” on page 107

Pitch Shifter options

Pitch Transpose Contains options that adjust pitch:

- **Semi-Tones** Transposes pitch in semi-tone increments, which equal musical half-notes (for example, the note C# is one semi-tone higher than C). A setting of 0 reflects the original pitch; +12 semi-tones is an octave higher; -12 semi-tones is an octave lower.
- **Cents** Adjusts pitch in fractions of semi-tones. Possible values range from -100 (one semi-tone lower) to +100 (one semi-tone higher).
- **Ratio** Determines the relationship between shifted and original frequency. Possible values range from 0.5 (an octave lower) to 2.0 (an octave higher).

Precision Determines sound quality, with the High setting taking longest to process. Use the Low setting for 8-bit or low-quality audio, and use the High setting for professionally recorded audio.



To quickly determine which Precision setting to use, process a small selected range at each setting until you find the best balance of quality and processing time.

Pitch Settings Control how audio is processed:

- **Splicing Frequency** Determines the size of each chunk of audio data. (The Pitch Shifter effect divides audio into very small chunks for processing.) The higher the value, the more precise the placement of stretched audio over time. However, artifacts become more noticeable as values go up. At higher Precision settings, a lower Splicing Frequency may add stutter or echo. If the frequency is too high, sound becomes tinny and voices have a tunnel-like quality.
- **Overlapping** Determines how much each chunk of audio data overlaps with the previous and next ones. If stretching produces a chorus effect, lower the Overlapping percentage. If doing so produces a choppy sound, adjust the percentage to strike a balance between choppiness and chorusing. Values range from 0 to 50%.
- **Use Appropriate Default Settings** Applies good default values for Splicing Frequency and Overlapping.

Stretch effect (Edit View only)

The Time And Pitch > 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 time.

See also

“About process effects” on page 104

“Apply individual effects in Edit View” on page 107

“Use effect presets” on page 104

Stretch options

Constant Stretch or Gliding Stretch Performs either 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% shrinks the wave). In Constant Stretch mode, you set a fixed Stretch percentage. In Gliding Stretch mode, you set Initial and Final percentages to stretch the waveform in a linear fashion from one ratio to another. Changes in the sliders are reflected in the Ratio and Length boxes.

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 $(\text{initial} + \text{final}) / 2$ 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 semitone (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. You can quickly process 8-bit or low-quality audio files 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 four 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 up.

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 raise the pitch, while percentages above 100 decrease the tempo and lower the pitch.
- **Constant Vowels** Preserves the sound of vowels in stretched vocals. This option requires substantial processing; try it on a small selection before applying it to a larger one.

Pitch And Time Settings (For Constant Stretch) Provide the following options:

- **Solo Instrument Or Voice** More accurately adjusts a solo performance.
- **Preserve Speech Characteristics** Maintains realism in speech.
- **Formant Shift** Determines how formants adjust to pitch shifts. The default value of zero adjusts formants together with pitch shifts, maintaining timbre and realism. Values above zero produce higher timbres (making a male voice sound female, for example). Values below zero do the reverse.

Pitch And Time Settings (For Gliding Stretch) 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.



In Low Precision mode, 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.

Multitrack effects

Envelope Follower effect (Multitrack View only)

The Multitrack > 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.

See also



“Select ranges in Multitrack View” on page 181

“About process effects” on page 104

“Control effects settings with graphs” on page 104

“Use effect presets” on page 104

Apply the Envelope Follower effect

- 1 In separate tracks, position two audio clips so that the sections you want to process together overlap.
- 2 Select the Hybrid tool  or the Time Selection tool .
- 3 In the Main panel, select the range you want to process.
- 4 Ctrl-click both audio clips.

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 Choose Effects > Envelope Follower, and set the desired options.

Envelope Follower 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** Is a mode 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).

Frequency Band Splitter effect (Multitrack View only)

The Multitrack > Frequency Band Splitter effect lets you take a selected audio 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).



See also

“Select ranges in Multitrack View” on page 181

“About process effects” on page 104

“Use effect presets” on page 104

Apply the Frequency Band Splitter effect

- 1 Select the Hybrid tool  or the Time Selection tool .
- 2 In the Main panel, select the range you want to process.
- 3 Ctrl-click the audio clip you want to process. (The Frequency Band Splitter effect is unavailable if multiple clips are selected.)

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.

- 4 Choose Effects > Frequency Band Splitter, and set the desired options.

Frequency Band Splitter 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 audio 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.

Vocoder effect (Multitrack View only)

The Multitrack > Vocoder effect simulates a hardware vocoder. 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).



See also

“Select ranges in Multitrack View” on page 181

“About process effects” on page 104

“Use effect presets” on page 104

Apply the Vocoder effect

- 1 In separate tracks, position two audio clips so that the sections you want to process together overlap.
- 2 Select the Hybrid tool  or the Time Selection tool .
- 3 In the Main panel, select the range you want to process.
- 4 Ctrl-click the audio 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 Choose Effects > Vocoder, and set the desired options.

Vocoder 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.*

Chapter 9: Mixing multitrack sessions

Multitrack View overview

About multitrack sessions

In Multitrack View, you can mix together multiple audio and MIDI tracks to create layered soundtracks and elaborate musical compositions. You can record and mix unlimited tracks, and each track can contain as many clips as you need—the only limits are hard disk space and processing power. When you're happy with a mix, you can export a mixdown file for use on CD, the web, and more.

Multitrack View is an extremely flexible, real-time editing environment, so you can change settings during playback and immediately hear the results. While listening to a session, for example, you can adjust track volume to properly blend tracks together. Any changes you make are impermanent, or *nondestructive*. If a mix doesn't sound good next week, or even next year, you can simply remix the original source files, freely applying and removing effects to create different sonic textures.

Adobe Audition saves information about source files and mix settings 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.

See also

“Create a new session” on page 46

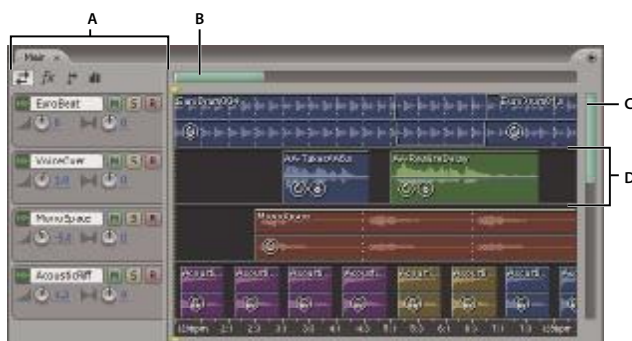
“Save multitrack sessions” on page 241

“Undo or redo changes” on page 101

“Comparing Edit View and Multitrack View” on page 20

Editing sessions in the Main panel and Mixer


In Multitrack View, the Main panel provides several elements that help you mix and edit sessions. In the track controls on the left, you adjust track-specific settings, such as volume and pan. In the timeline on the right, you edit the clips and automation envelopes in each track.




Main (Timeline) panel in Multitrack View

A. Track controls B. Horizontal scroll bar C. Vertical scroll bar D. Track

The Mixer (Window > Mixer) provides an alternative view of a session, revealing many more tracks and controls simultaneously, without showing clips. The Mixer is ideal for recording automation or mixing large sessions with many tracks.

 To quickly toggle between the Main panel and Mixer, dock them in the same panel group. (See “Dock, group, or float panels” on page 23.)



Controls in the Mixer:
 A. Inputs B. Effects C. Sends D. Equalization E. Volume F. Outputs G. Automation

See also

“Add, insert, or delete tracks” on page 182

“About clips” on page 191



“Automating track settings” on page 204

Scroll to the same tracks in the Main panel and Mixer

By default, Adobe Audition scrolls to different tracks in the Main panel and Mixer, maintaining unique positions in each. To automatically scroll to the same tracks, link the Main panel and Mixer.



❖ From the Options menu, select Link Mixer And Timeline.

Select ranges in Multitrack View

To select ranges in Multitrack View, you 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 and move clips with separate tools, use the Time Selection and Move/Copy Clip tools. (See “About clips” on page 191.)



Simultaneously selecting a range and clips in the Main panel (Upper two clips are selected, third isn't)

- 1 In the toolbar, select either the Hybrid tool  or the Time Selection tool .
- 2 In the Main panel, 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 a marquee.


Insert or delete time in a session

You can insert silence into or delete a selected range from a session.

- 1 Place the start-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 start-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 frozen tracks to insert or delete time in a session. (See “Freeze effects and improve performance” on page 109.)

Basic track controls

Understanding different track types

Multitrack sessions can include five different types of tracks:

- Audio tracks  contain either imported audio or clips recorded in the current session. These tracks offer the widest range of controls, letting you specify inputs and outputs, apply effects and equalization, route audio to sends and buses, and automate mixes.
- MIDI tracks  let you access a sequencer in which you can import, record, and edit MIDI compositions. Adobe Audition uses VSTi-based virtual instruments to automatically convert the MIDI data into audio, providing MIDI tracks with nearly all the controls that audio tracks offer.
- Video tracks  contain an imported video clip. A session can include a maximum of one such clip at a time. Because this clip exists primarily for visual reference, track controls are limited to thumbnail display options. However, you can watch a preview in the Video panel (Window > Video).
- Bus tracks  let you combine the outputs of several audio tracks or sends and control them collectively.
- The Master track , which each session always contains, lets you easily combine the outputs of multiple tracks and buses and control them with a single fader.

See also

“Applying effects in Multitrack View” on page 107

“Track routing and EQ controls” on page 185

“MIDI sequencing basics” on page 213

“Importing video and working with video clips” on page 228

Add, insert, or delete tracks

If you regularly create sessions with a specific track structure, use the Add Tracks command to instantly add multiple tracks of various types. Or, individually insert new tracks to precisely place them.

Note: A multitrack session supports only one video track, which Adobe Audition always inserts at the top of the Main panel.

❖ In the Main panel or Mixer, do any of the following:

- To add multiple tracks of various types, choose Insert > Add Tracks, and specify the number and placement of new audio, MIDI, and bus tracks. Select Add Video Track if needed, and click OK.
- To insert a track, select the track you want to precede it, and then choose Insert > [type of] Track.
- To delete a track, select it, and choose Edit > Delete Selected Track.

See also

“Insert an audio file into a session” on page 47

“Insert a video file into a session” on page 228

“Insert MIDI into a session and access the Sequencer” on page 214

Name or move tracks

You can name tracks to better identify them, or move them to display related tracks together.

- In the Main panel or Mixer, type in the name text box.



Name text box in the Main panel

- Position the pointer to the left of the track name, and then drag either up or down in the Main panel, or right or left in the Mixer.



Moving a track in the Main panel

Vertically zoom an individual track

When you use the vertical options in the Zoom panel, all tracks zoom simultaneously. If a session contains many tracks, however, you may prefer to zoom them individually.

- ❖ In the track controls, drag the top or bottom border of the track up or down.

💡 To horizontally resize all track controls, drag the right border.



Vertically zooming an individual track


See also

“Zooming audio” on page 20

“Show or hide track routing and EQ controls” on page 185

Set track output volume

❖ Do any of the following:

- In the Main panel, drag the Volume knob . Hold down Shift to change settings in 10-unit increments. Hold down Ctrl to change settings in fractional increments of one-tenth.
- In the Mixer, drag the track fader, click above or below to move incrementally, or Alt-click above or below to move to the click point.



Alt-click knobs and faders to return them to zero (unity gain).

See also



“Assign audio inputs and outputs to tracks” on page 186

“Pan tracks in stereo or sum them to mono” on page 189

“Automating track settings” on page 204

Mute and solo 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 mute a track, click its Mute button  in the Main panel or Mixer.
- To solo a track, click its Solo button  in the Main panel or Mixer. (Ctrl-click to automatically remove other tracks from Solo mode.)

Note: *The Ctrl-click shortcut doesn't apply to bus tracks; when you solo a bus, assigned tracks are always placed in Solo mode.*

See also

“Record audio clips in Multitrack View” on page 57

Apply an identical setting to all tracks

To increase your efficiency, you can quickly apply several settings to an entire session.

- ❖ Hold down Ctrl+Shift, and then select a Solo, Mute, Record, Input, or Output setting for any track.

See also

“Mute and solo tracks” on page 184

“Record audio clips in Multitrack View” on page 57

“Assign audio inputs and outputs to tracks” on page 186

Duplicate a track

To perfectly copy all clips, effects, equalization, and envelopes in a track, duplicate it. Duplicate tracks provide a great starting point for new adjustments, helping you compare different processing and automation settings.

- 1 In the Main panel or Mixer, select a track.
- 2 Choose Edit > Duplicate Selected Track.

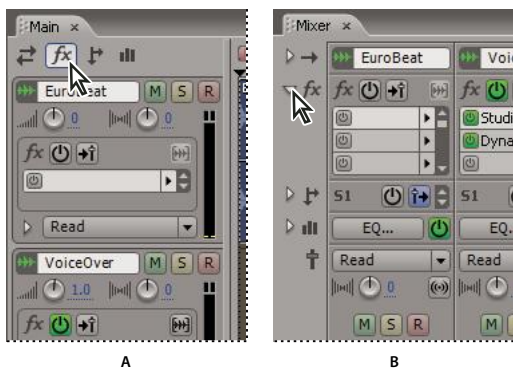
See also

- “Mute and solo tracks” on page 184
- “Applying effects in Multitrack View” on page 107
- “Automating mixes” on page 203

Track routing and EQ controls





Show or hide track routing and EQ controls

Although the wide variety of routing and EQ controls may seem intimidating at first, the controls for each track are identical, so once you’ve learned one, you’ve learned them all.



Showing and hiding sets of audio track controls
 A. Main panel B. Mixer

❖ Do either of the following:

- On the left side of the Mixer, click the Show/Hide triangle for one or more sets of controls.
- In the upper left corner of the Main panel, click the button for Inputs/Outputs , Effects , Sends , or EQ .



In the Main panel, drag the right or bottom border in the track controls to show more or less detail.



Dragging the right border of the track controls to reveal more or less detail


See also

- “Vertically zoom an individual track” on page 183

Assign audio inputs and outputs to tracks

❖ In the Inputs/Outputs area  of the Main panel or Mixer, do the following:

- From the Input menu, choose a hardware input.
- From the Output menu, choose a bus, the Master track, or a hardware output.

 *The list of available hardware ports is determined by settings in the Audio Hardware Setup dialog box. (See “Set audio inputs and outputs” on page 31.)*

See also

“Connect to MIDI inputs and outputs” on page 215


“Record audio clips in Multitrack View” on page 57


“Set track output volume” on page 184

“Routing audio to buses, sends, and the Master track” on page 186

Boost or lower input gain

To best optimize input levels, set appropriate sound card levels during recording. In some mixes, however, a previously recorded track may be so loud or quiet that you can't fully compensate with track volume settings. To address this situation, boost or lower input gain. (When boosting, be careful to avoid distortion.)


❖ In the Input Controls area  of the Mixer, drag the Input Gain knob.

 *Hold down Shift to change settings in 10-unit increments. Hold down Ctrl to change settings in fractional increments of one-tenth.*

Invert the phase of an input


If a pair of stereo inputs such as overhead drum microphones are out-of-phase, you'll hear duller sound and a narrower stereo image. To correct the problem, invert the phase of one the inputs.


❖ In the Input Controls area  of the Mixer, click the Phase Invert button .

 *For a definition of this key audio concept, see “phase” on page 277.*

Routing audio to buses, sends, and the Master track

Buses, sends, and the Master track let you route multiple track outputs to one set of controls. With these combined controls, you can efficiently organize and mix a session.

 *To save the output of an audio, bus, or Master track to a file, see “Export a session to an audio file” on page 242.*




An example of audio routing for tracks
 A. Vocal B. Reverb bus receiving vocal and guitar sends C. Guitar D. Drums bus combining drum outputs E. Bass outputting direct to hardware F. Master track G. Hardware outputs

Understanding bus tracks

With bus tracks, you can combine the outputs of several audio tracks or sends and control them collectively. For example, to control the volume of multiple drum tracks with a single fader, output all the tracks to one bus. Or, to optimize system performance, apply a single reverb effect to a bus track, and then output sends from multiple tracks to that bus. (Individually applying the same reverb to multiple tracks would inefficiently use CPU resources.)


Although bus tracks lack a hardware input, they have all the other features of audio tracks. You can apply effects and equalization and automate your changes over time. For most mixes, you'll output buses to hardware ports or the Master track. If you need to combine buses, however, you can even output them to other buses.



An example of bus routing:
 A. Drum kit bus B. Hand drum bus C. Combined drums bus outputting to either the Master track or hardware

Understanding sends

Sends let you route audio from a track to multiple buses, creating tremendous signal-routing flexibility. Each track provides up to 16 sends, which you configure independently from the track output. For example, you can output an unprocessed track directly to a hardware port, but output Send 1 to a reverb bus and Send 2 to a headphone bus. (A headphone bus lets performers hear a unique mix during recording. Drummers, for example, may prefer a louder bass track.)



Sending tracks to multiple buses

A. Send 1 outputs to delay bus B. Send 2 outputs to reverb bus C. Master track combines vocal, guitar, delay, and reverb outputs

Understanding the Master track

A session always contains one Master track, so you can easily combine the outputs of multiple tracks and buses and control them with a single fader. Because the Master track exists at the very end of the signal path, it offers fewer routing options than audio and bus tracks. The Master track can't directly connect to audio inputs, or output to sends or buses; it can only output directly to hardware ports.




A B

The Master track always exists at the end of the signal path.

A. Main panel B. Mixer

Set up a send

When you set up a send, you determine the volume and stereo pan it outputs to an assigned bus. You also place the send either pre- or post-fader. Pre-fader sends aren't affected by track volume; post-fader sends are. (For example, if you output a pre-fader send to a reverb bus, the reverb continues after you fade out dry audio. If you instead output a post-fader send, the reverb fades out in unison with dry audio.)



Pre- and post-fader effect and send routing for each track
 A. Input gain B. EQ C. Track volume D. Track mute E. Send F. Effects Rack

- 1 In the Sends area of the Main panel or Mixer, click the Send Power button .
- 2 Click the Pre-Fader/Post-Fader button to place the send either before track volume or after .
- 3 Set Volume and Pan.
- 4 From the Output menu, select a bus.

See also

- “Routing audio to buses, sends, and the Master track” on page 186
- “Set track output volume” on page 184
- “Pan tracks in stereo or sum them to mono” on page 189
- “Insert effects before or after sends and EQ” on page 109

Pan tracks in stereo or sum them to mono

- Drag the Pan knob in the Main panel or Mixer.

Hold down Shift to change settings in 10-unit increments. Hold down Ctrl to change settings in fractional increments of one-tenth.

- Click the Sum To Mono button in the Mixer.

See also

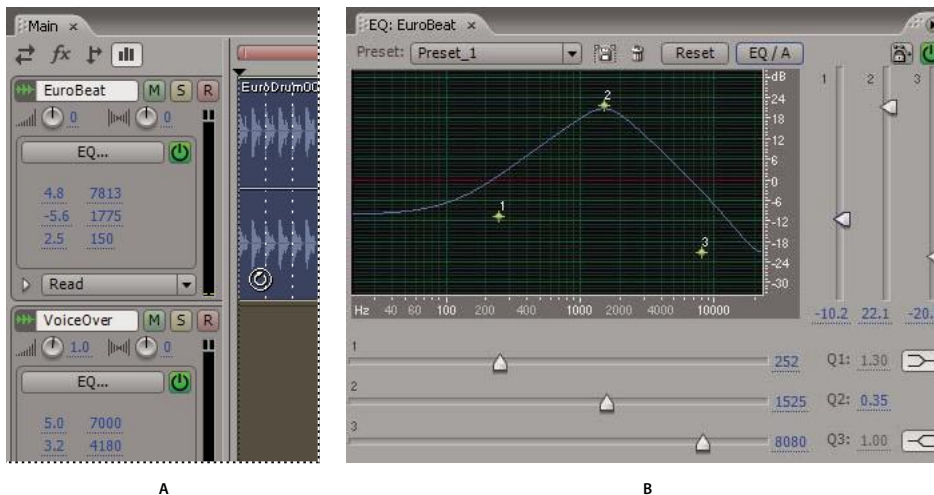
- “Set track output volume” on page 184
- “Automating track settings” on page 204

Equalizing tracks

For each track, Multitrack View provides a three-band parametric equalizer that you can access in the Main, Mixer, or Track EQ panel. The Main and Mixer panels provide quick access to commonly used settings; the Track EQ panel provides more precise and sophisticated controls.

In the EQ area of the Main panel and Mixer, the blue text shows the current equalization for three frequency bands. You drag across this text to change center frequency and adjust boost or attenuation.

In the Track EQ panel, you use a graph to visually adjust equalization settings. In addition to options offered by the Main panel and Mixer, the Track EQ panel provides access to presets, bandwidth controls, and other advanced options.





Equalizing tracks
A. Numeric controls in the Main panel B. Graphical controls in the Track EQ panel

See also

“Use effect presets” on page 104

Equalize a track in the Main panel or Mixer


- 1 In the EQ area , click the EQ Power button .
- 2 Drag across the blue text for each band to specify boost or attenuation in decibels, and center frequency in Hertz.
- 3 (Optional) Click the EQ button to open the Track EQ panel and see a graphical view of your settings.

Equalize a track in the Track EQ panel

- 1 Select the track, and then choose Window > Track EQ.
- 2 Set the following options:

Preset options Let you store and apply identical settings to multiple tracks.

EQ/A-B Lets you compare two different equalization settings. 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.

Track EQ Safe During Write  Protects all settings from inadvertent changes that can occur when recording automation.

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. Initially, the leftmost control point represents the low frequency, and the middle and rightmost control points correspond to the middle and high frequencies.

Vertical amplitude sliders Boost or attenuate the three frequency bands. Either drag the sliders or click the blue text to enter precise values.

Horizontal frequency sliders Determine the center frequency for each band. Either drag the sliders or click the blue text to enter precise values.




Q values Define the width of each band. (Higher values equal narrower bandwidth, and vice versa.) Either drag across the text or click it to enter precise values.

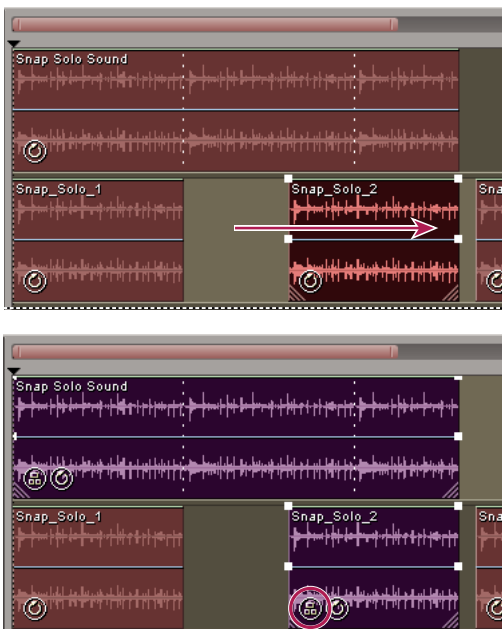
Band/Low Shelf, Band/High Shelf buttons Switch the first and third bands from shelving to peaking, and vice versa.

Arranging clips

About clips

When you insert an audio 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 Main panel, 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 those tasks.



Aligning and grouping two clips


See also


“Select ranges in Multitrack View” on page 181


Select and move clips

❖ Do any of the following:


- To select an individual clip, click the clip in the Main panel.
- To select all clips in a track, select the track, and then 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.
- To move selected clips, select the Move/Copy Clip tool  in the toolbar, and then drag the clips.

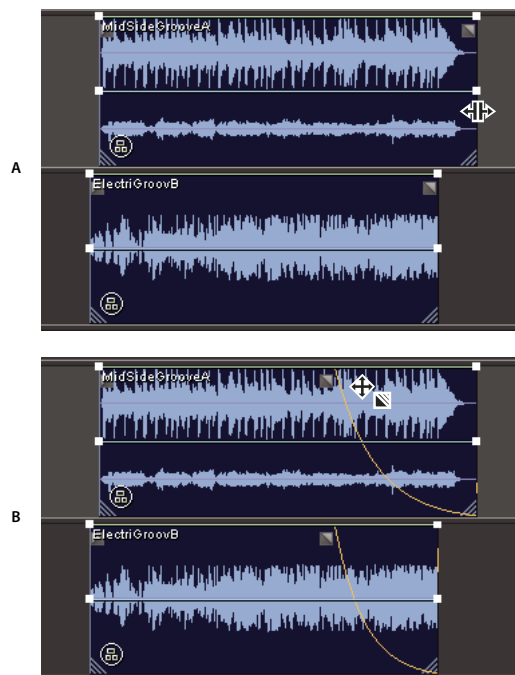
 If you prefer, select the Hybrid tool, and then right-click and drag the clips.

Grouping clips

Group clips to more efficiently organize, edit, and mix a session. For example, group related guitar clips together to easily identify, select, and move them. Grouped clips appear with the group icon  and share a common color.

You can quickly edit grouped clips as a unit. To trim a group, simply drag its outermost edges. To fade a group, drag an on-clip icon.

 Changes to clip mute and lock properties affect all audio clips in a group. (See “Set audio clip properties” on page 202.)



Collectively editing grouped clips

A. Drag outermost edge to trim group B. Drag on-clip icon to fade group

See also

“Trimming and extending clips” on page 195

“Fade or crossfade clips in a track” on page 197

Group clips

- 1 Hold down the Ctrl key, and click each clip you want in the group.
- 2 Choose Clip > Group Clips. Alternatively, right-click any clip in the group, and choose Group Clips.

Ungroup clips

- ❖ Select any clip in the group, and choose Clip > Group Clips.

Change the color for a group

- 1 Select any clip in the group, and choose Clip > Group Color.
- 2 Select a color, and click OK.

Align clips

You can align the left or right edges of multiple clips, giving them the same start or end point.

- 1 Hold down Ctrl, and select the clips.
- 2 Choose Clip > Align Left or Clip > Align Right.

Note: Because the relative position of grouped clips is fixed, you must ungroup them to align them.

Snap to clip and loop endpoints

Snapping lets you quickly align clips with loops and other clips. If snapping is enabled, both dragged clips and the start-time indicator snap to loop endpoints and clip edges.

While you drag a clip, a white line appears in the Main panel 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.

- ❖ 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.




This section describes snapping options that are unique to Multitrack View. For information about snapping options that Multitrack View shares with Edit View, see “Snap to markers, rulers, frames, and zero crossings” on page 76.


Copy a clip

You can create two types of copied audio clips: reference copies that share source files and unique copies that have independent source files. 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 a 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 quickly copy a reference, press **Ctrl+C**.

- 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.

See also

“Repeat a clip” on page 194

Repeat a clip


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 “Copy a clip” on page 193.


- 1 Select the clip, and choose **Clip > Duplicate**.
 - 2 Set the following options, and then click **OK**:
- 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 221.

- **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 during recording and playback” on page 54.) 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**.

Insert 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 “Punch into a selected range in Multitrack View” on page 58.)

- 1 Select a range in the Main panel.

2 Choose Insert > Empty Audio Clip, and then choose one of the following:

- In Selected Track (stereo)
- In Selected Track (mono)

Display, remove, or play hidden clips

If clips overlap, you can display, remove, or play hidden clips. (By default, Adobe Audition plays only displayed clips.)

- To display individual hidden clips, select the overlapping clip, and choose Clip > Bring To Front > [clip name].
- To display hidden clips throughout a session, choose Edit > Check for Hidden Clips.
- To remove hidden clips, select the overlapping clip, and choose Clip > Remove Hidden Clips.
- To play hidden clips, select the overlapping clip, and choose Clip > Play Hidden Clips.

Remove selected clips

You can remove selected clips from a session but keep their source files available in the Insert menu, the Files panel, and Edit View. Alternatively, you can destroy selected clips to both remove them from a session and close their source files.

- To remove selected clips, choose Clip > Remove, or press Delete.
- To destroy selected clips, choose Clip > Destroy.

Editing clips

Trimming and extending clips

You can trim or extend audio 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.

See also



“Comparing Edit View and Multitrack View” on page 20

“Extend or shorten a looped clip” on page 222

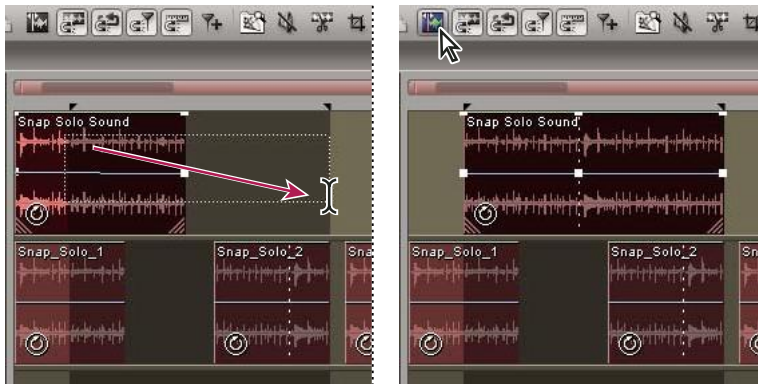
“Grouping clips” on page 192

“Display the shortcut bar” on page 27

Trim or extend clips by selecting a range



- 1 In the toolbar, click the Time Selection tool  or the Hybrid tool .
- 2 Drag across one or more clips to select them and a range.
- 3 Do one of the following:
 - To trim clips to the range, choose Clip > Trim.
 - To remove the range from clips and leave a gap in the timeline, choose Edit > Delete.
 - To remove the range and collapse the gap in the timeline, choose Edit > Ripple Delete.

- To adjust clip edges to the range, choose Clip > Adjust Boundaries. (To reveal more of a previously edited clip, extend the range beyond the current clip edges.)



Using the Adjust Boundaries command to reveal more of a previously edited clip

Trim or extend clips by dragging



- 1 Choose View > Enable Clip Edge Dragging.
- 2 In the Main panel, position the cursor over the left or right edge of the clip. The edge-dragging icon  appears. (If the time stretch icon  appears instead, position the cursor above the corner handle.)
- 3 Drag clip edges.

Shift the contents of a trimmed or looped clip

You can slip edit a trimmed or looped clip to shift its contents within clip edges.




Shift clip contents within clip edges

- 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.

Return to the full, original version of a clip

- ❖ Select the clip, and choose Clip > Full.


 The Full command doesn't apply to loops; to extend or shorten a loop, simply drag its edges.

Permanently edit a clip's source file in Edit View

- ❖ Double-click the clip.



Splitting and rejoining clips

The Split command functions like a traditional tape splice; it cuts audio 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

Split a clip

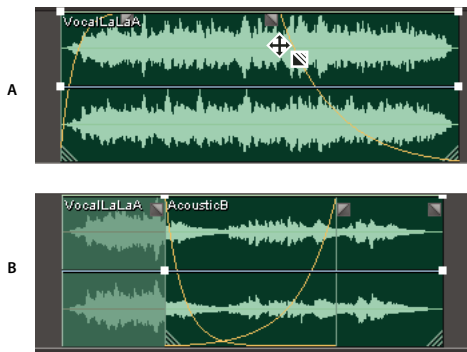
- 1 In the toolbar, click the Time Selection tool  or the Hybrid tool .
- 2 Do either of the following:
 - To split the clip in two, click where you want the split to occur.
 - To split the clip into three, drag across it to specify two split points (one at the beginning of the selection; one at the end).
- 3 Choose Clip > Split.

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 Select one of the clips, and choose Clip > Merge/Rejoin Split.

Fade or crossfade clips in a track

On-clip fade and crossfade controls let you visually adjust fade curves and duration. Controls for fade ins and fade outs always appear in the upper-left and upper-right corners of clips. Controls for crossfades appear only when you overlap clips.





On-clip controls
A. Drag controls in clip corners to fade in and out B. Overlap clips to crossfade

See also

“Select and move clips” on page 192



“Select ranges in Multitrack View” on page 181

Fade a clip in or out

❖ In the upper-left or upper-right corner of the clip, drag the fade icon   inward to determine fade length, and drag up or down to adjust the fade curve.

Crossfade overlapping clips

When you crossfade clips on the same track, you overlap them to determine the size of the transition region (the larger the overlapping area, the longer the transition).

- 1 Place two clips on the same track.
- 2 Move the clips so they overlap.
- 3 At the top of the overlapping area, drag the left  or right  fade icon up or down to adjust the fade curves.



If you want to apply unique effects to each clip in a crossfade, see “Crossfade clips in separate tracks” on page 199.

Fade options

To access the following options for fades and overlapping crossfades, select a clip, and then either right-click a fade icon in the Main panel, or choose Clip > On-Clip Fades.

Linear or Cosine Applies either an even, linear fade or an S-shaped fade that starts slowly, then rapidly changes amplitude, and ends slowly.

Crossfade Vertical Adjustments Determine how the left and right fade curves interact when you drag them up and down:

- **Unlinked** Lets you independently adjust each curve.
- **Linked** Lets you boost or attenuate one curve, and automatically does the reverse with the other.
- **Linked (Symmetrical)** Lets you identically adjust both curves.

Remove Deletes the fade or crossfade.


Set Fade In (or Out) As Default Fade Curve Saves the current fade curve as the default for Multitrack View.

Allow Vertical Fade Adjustments Enables you to adjust fade curves. Disable this option if you're inadvertently adjusting these curves when attempting to perform other tasks, such as editing volume or pan envelopes.

Automatically Crossfade Crossfades overlapping clips. Deselect this option if automatic crossfades are undesirable or interfere with other tasks, such as trimming clips.

Crossfade clips in separate tracks

When you crossfade clips in separate tracks, you select a range that determines the transition region. To create a smooth transition, select a range that starts before the end of the first clip and extends beyond the beginning of the second clip.



Selecting range and two clips, and applying linear crossfade


- 1 Place two clips in 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 range for the crossfade.
- 4 Ctrl-click both clips.
- 5 Choose Clip > Fade Envelope Across Selection, and then choose one of the following:

Linear Produces an even crossfade.

Sinusoidal Produces a crossfade with a curved, sine-like slope.

Logarithmic In Fades in logarithmically, producing a steeper slope at the end of the fade.

Logarithmic Out Fades out logarithmically, producing a steeper slope at the beginning of the fade.

 The resulting fade curves are volume envelopes, which you can edit. (See “Edit automation envelopes” on page 207.)

See also

“Select and move clips” on page 192

“Select ranges in Multitrack View” on page 181

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 can 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

Time stretch a clip by dragging

- 1 Choose View > Enable Clip Time Stretching.
- 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.

Set specific time stretch properties

- 1 Select the clip, and choose Clip > 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. Solo Voice/Instrument and Preserve Formant options maintain realism.

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.

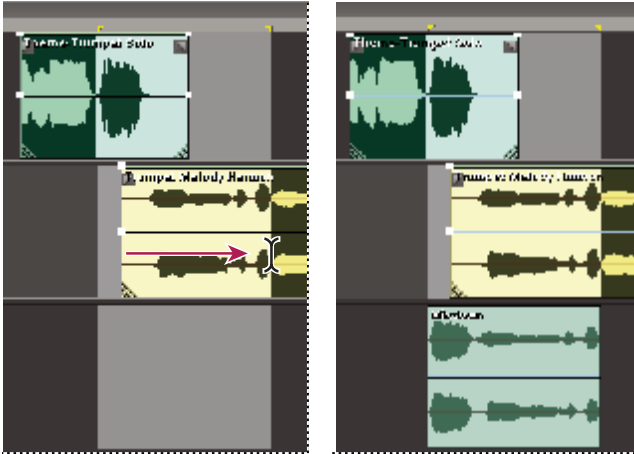
Hybrid Uses the current Time-Scale Stretch settings when you shorten the clip, and uses the current Beat Splice settings when you lengthen it.

Disable time stretching

- 1 Select a time-stretched clip, and choose Clip > Clip Time Stretch Properties.
- 2 Deselect Enable Time Stretching.

Create a single audio clip from multiple clips

You can combine the contents of multiple clips in the same time range, creating a single clip that you can quickly edit in either Multitrack View or Edit View.



Creating single clip from multiple clips in Multitrack View

- 1 In the Main panel, do any of the following:
 - Select a specific range to mix down. (See “Select ranges in Multitrack View” on page 181.)
 - Select specific clips to mix down.
 - Select nothing to mix down an entire session.
- 2 To combine the contents of the original clips, do either of the following:
 - To create a track and clip in Multitrack View, choose Edit > Bounce To New Track.
 - To create a file in Edit View, choose Edit > Mixdown To New File.
- 3 Choose a stereo or mono mixdown option.

See also

“Export a session to an audio file” on page 242

Create an audio clip from a bus, master, MIDI, or ReWire track

If you want to edit audio from a bus, master, MIDI, or ReWire track, create a clip from the track.

- 1 In the Main panel, select a time range for the new clip. (See “Select ranges in Multitrack View” on page 181.)
- 2 In the timeline, right-click a bus, master, MIDI, or ReWire track, choose Bounce *[track type]* To Track, and then choose Stereo or Mono.

Adobe Audition creates a new track with a clip that reflects the bus, master, MIDI, or ReWire mix.

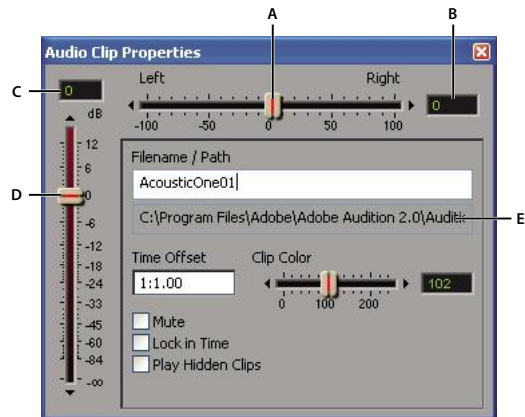
See also

“Export a session to an audio file” on page 242

“Synchronizing with ReWire” on page 208

Set 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.




Audio Clip Properties window

A. Pan slider B. Pan text box C. Volume text box D. Volume slider E. Pathname for source file

1 Select an audio clip, and choose Clip > 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.

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.

- To silence the clip, select Mute.
- To move the clip to a specific 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.)

See also

“Display, remove, or play hidden clips” on page 195

Automating mixes

Techniques for automating mixes

By automating mixes, you can change mix settings over time. For example, you can automatically increase volume during a critical musical passage and later reduce the volume in a gradual fade out.


Adobe Audition provides several techniques for automating mixes. You may find that one method best suits your working style and audio content. Typically, however, you'll probably use a combination of these techniques:

- To manually automate clip volume and pan settings, use clip envelopes.
- To manually automate track volume, pan, and effects settings, use track envelopes.
- To dynamically automate track settings while mixing, record track automation.

About automation envelopes

Automation envelopes visually indicate settings at specific points in time, and you can edit them by dragging edit points on envelope lines. With volume envelopes, for example, the line represents full volume at the top, and zero volume at the bottom.

Envelopes are nondestructive, so they don't change audio files in any way. If you open a file in Edit View, for example, you don't hear the effect of any envelopes applied in Multitrack View.



Clip and track envelopes in the Main panel
A. Clip envelope B. Track envelope


See also

“Automating clip settings” on page 203

“Automating track settings” on page 204

Automating clip settings

With clip envelopes, you can automate clip volume and pan settings. You can identify clip Volume and Pan envelopes by color and initial position. Volume envelopes are green lines initially placed across the top of clips. Pan envelopes are blue lines initially placed in the center. (With pan envelopes, the top of a clip represents full left, while the bottom represents full right.)



Two clip envelopes
 A. Volume envelope B. Pan envelope

See also

“Edit automation envelopes” on page 207

“About spline curves for graphs” on page 105

Show or hide clip envelopes

❖ From the View menu, choose either of the following:

- Show Clip Volume Envelopes
- Show Clip Pan Envelopes

Use spline curves for clip envelopes

❖ Right-click the clip containing the envelope, and choose Clip Envelopes > [envelope type] > Use Splines.

Rescale a clip volume envelope

If an envelope is too high or low, preventing you from raising or lowering edit points, you can rescale it.

- 1 Right-click the clip containing the envelope, and choose Clip Envelopes > Rescale Clip 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 properties for sessions” on page 210.)

Automating track settings

With track envelopes, you can change volume, pan, and effect settings over time. Adobe Audition displays track envelopes in automation lanes below each track. Each automated parameter has its own automation lane and envelope, which you edit just like clip envelopes.



Automating track settings in the Main panel
 A. Automation lane B. Envelope for parameter

See also

“Apply effects in Multitrack View” on page 107

“Edit automation envelopes” on page 207

“Editing sessions in the Main panel and Mixer” on page 179



“External Controllers preferences” on page 40

Manually create track envelopes

While a session is playing or stopped, you can manually create track envelopes, precisely changing track settings at specific points in time.



Show/Hide Automation Lanes button in Main panel

- 1 In Main panel, click the Show/Hide Automation Lanes button  for the track you want to automate.
- 2 From the pop-up menu in the automation lane, choose a parameter to automate.
- 3 In the created lane, click and drag the envelope to add and adjust edit points.
- 4 (Optional) Click the Show Additional Automation Lane button  to add another lane, and then repeat steps 2 and 3.

Record track automation while mixing

While playing a session, you can record adjustments you make to track volume, pan, and effect settings, creating a mix that dynamically evolves over time. Adobe Audition automatically converts your adjustments into track envelopes, which you can edit with precision.



With an external controller such as the Mackie Control, you can adjust multiple settings simultaneously.

- 1 In the Main panel, position the start-time indicator where you want to start recording automation.
- 2 Choose an option from the Track Automation Mode menu.
- 3 To start recording automation, press the spacebar, or click Play in the Transport panel. As audio plays, adjust track or effect settings in the Main, Mixer, or Effects Rack panels.
- 4 To stop recording automation, press the spacebar, or click Stop in the Transport panel.



To prevent recorded automation from creating excessive or irregular edit points, optimize automation settings in the Preferences dialog box. (See “Multitrack preferences” on page 38.)

Track Automation Mode options

When recording automation in the Main panel or Mixer, you can choose one of the following modes for each track:



Track Automation Mode menu

Off Ignores track envelopes during playback and mixdown, but continues to display envelopes so you can manually add or adjust edit points.

Read Applies track envelopes during playback and mixdown, but doesn't record any changes you make to them. (You can preview such changes, but edit points return to recorded settings.)



Latch Records adjustments you make to settings and creates corresponding edit points on track envelopes. Begins recording when you first adjust a setting, and continues to record new settings until playback stops.

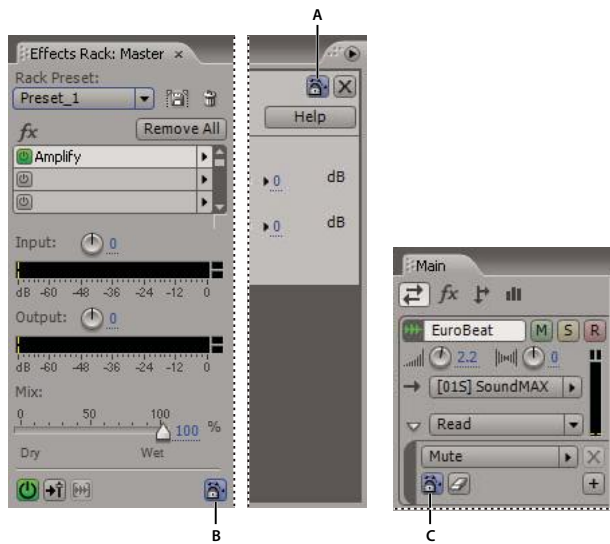
Touch Similar to Latch, but returns settings to previously recorded values when you stop adjusting them.

Write Similar to Latch, but records current settings as edit points when playback starts, without waiting for a setting to change.

Protect settings when recording automation

When recording automation, you may want to protect certain settings from inadvertent changes. In the Effects Rack, you can protect entire racks or individual effects; in the Main panel, you can protect individual parameters.


- In the Effects Rack, click the Safe During Write button for either the entire rack or individual effects.
- In the Main panel, click the Safe During Write button  for individual parameters. (To reveal a parameter, click the Show/Hide Automation Lanes button  in the track controls. Then choose the parameter from the pop-up menu in the automation lane.)




Safe During Write buttons in the Effects Rack and Main panel:
A. Individual effect B. Entire rack C. Individual parameter

Edit automation envelopes

After you create clip or track envelopes, you can fine-tune them by dragging edit points while a session is playing or stopped.



Adjusting edit point for envelope

- 1 (Clip envelopes only) Click the Edit Clip Envelopes button  in the shortcut bar. Then select the clip.
- 2 Do any of the following:
 - To add an edit point, click the envelope.
 - To remove an edit point, drag it off the clip or track.
 - To move an edit point, drag it. (To maintain time position, hold down Shift while dragging.)
 - To move all edit points up or down by the same percentage, hold down Ctrl while you drag.
 - To move all edit 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.)





To move a group of points in a track envelope, select a range, and then Ctrl- or Alt-drag any point in the range.

Remove edit points from an envelope

- 1 (Optional) To remove edit points from a specific time range, select that range in the Main panel.

2 Do either of the following:

- For a clip envelope, right-click the clip, and choose Clip Envelopes > [envelope type] > Clear Selected Points.
- For a track envelope, show the desired automation lane (click the Show/Hide Automation Lanes button  in the track controls, and choose the specific parameter from the pop-up menu). Then click the Clear Edit Points button .

Synchronizing sessions with ReWire and SMPTE

Synchronizing with ReWire

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*.


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 Audio Hardware Setup dialog box. (See “Set audio inputs and outputs” on page 31.) You can also, however, preview individual modules in the slave application to hear them independently from the Adobe Audition session. When you do, the modules send audio through the sound card specified in the Sounds and Audio Devices control panel.

Synchronize a ReWire slave with a session

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

- 1 In Multitrack View, choose Edit > ReWire Setup, and click 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 track Input menus.
- 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 In the ReWire slave application, open a session with audio you want to access in Adobe Audition.

Adobe Audition serves as the exclusive ReWire host until you close the application or manually disable ReWire.

 If you notice a timeline offset between Adobe Audition and the slave application, lower the buffer size for your sound card (choose Edit > Audio Hardware Setup, click the Multitrack View tab, and click Control Panel). Because extremely low buffer sizes may cause audio to drop out, you may need to try different values to find one that is acceptable.

Manually disable ReWire


Only one ReWire host can be active at a time. If you want another application to serve as the host, you must either close Adobe Audition or manually disable ReWire.

- 1 In Multitrack View, choose Edit > ReWire Setup.
- 2 Click the ReWire tab, and deselect all active slave applications.
- 3 Click Disable, and then click OK.

Assign a ReWire device to an audio track

After you synchronize ReWire applications with a session, you can assign ReWire devices to audio tracks and set volume, pan, effects, and routing options. Note, however, that saved sessions store only settings in Adobe Audition; if you change settings in the ReWire slave application, you must save them separately.

- ❖ In the track controls, click the Input menu, and choose ReWire > [device].

 To the right of the track controls, a blank clip representing ReWire audio data appears. If you want to edit this data, convert it to an audio clip. (See “Create an audio clip from a bus, master, MIDI, or ReWire track” on page 201.)



ReWire clips are blank in the Main panel

Synchronizing with SMPTE

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. 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.

Use a session as a SMPTE master

As a SMPTE master, Adobe Audition sends timecode through a MIDI Out port of your system.

- 1 Choose Edit > Preferences, and click the SMPTE/MTC tab.
- 2 Choose a port from the MIDI Output Device menu, and click OK. (See “SMPTE/MTC preferences” on page 40.)
- 3 Choose Options > SMPTE/MTC Master Enable.
- 4 Select the desired SMPTE time format for the timeline. (See “Monitoring time during recording and playback” on page 54.)

Use a session as a SMPTE slave

As a SMPTE slave, Adobe Audition receives timecode from a MIDI In port of your system.

- 1 Choose Edit > Preferences, and click the SMPTE/MTC tab.
- 2 Choose a port from the MIDI Input Device menu, set Synchronization options, and then click OK. (See “SMPTE/MTC preferences” on page 40.)
- 3 Choose Options > SMPTE Start Offset, click Format, and select the desired SMPTE time format.
- 4 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.

5 Choose Options > SMPTE/MTC Slave Enable.

In the left of the status bar, Adobe Audition reports the following synchronization statuses:

- 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.

Synchronize to SMPTE with sample accuracy


If your MIDI interface supports sample-accurate synchronization, you can more precisely lock a session to timecode.

❖ Choose Options > Sample Accurate Sync.

Setting advanced session properties

Set 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 panel. See “Set tempo, time signature, and key for a session” on page 223.

1 Choose View > Advanced Session Properties.

2 Set options as desired, and click OK.

General properties 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 during recording and playback” on page 54.

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 properties for sessions

In the Advanced Session Properties dialog box, click the Mixing tab to set the following session-specific options:

 To set global properties for sessions, use the Multitrack tab in the Preferences dialog box. (See “Multitrack preferences” on page 38.)

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.

Clip 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.

Tempo properties for sessions

In the Advanced Session Properties dialog box, click the Tempo tab to 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 lets you line up the metronome with an existing song that you're adding more to.
- 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 properties for sessions

In the Advanced Session Properties dialog box, click the Metronome tab to 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.

Output To Sets the sound card port for the metronome.

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.)

Delete Removes the current Signature option.

Pattern Shows the current metronome pattern.

Add notes to sessions

Session notes can help you recall details about a session or communicate those details to someone else. For example, notes could list the musicians for the session, indicate the instrument each played, and detail alternative EQ settings for those instruments.

- ❖ Choose View > Advanced Session Properties, and click Notes.

Chapter 10: Composing with MIDI

MIDI sequencing basics

Understanding MIDI data and VST instruments

Digital audio and MIDI are fundamentally different: Digital audio represents a sound wave; MIDI provides instructions for musical instruments. Like a tape recorder, a digital audio file tries to exactly reflect an audio event, whether it's a musical performance, a person talking, or any other sound. MIDI instructions, on the other hand, are like sheet music that lets digital instruments re-create a composition.

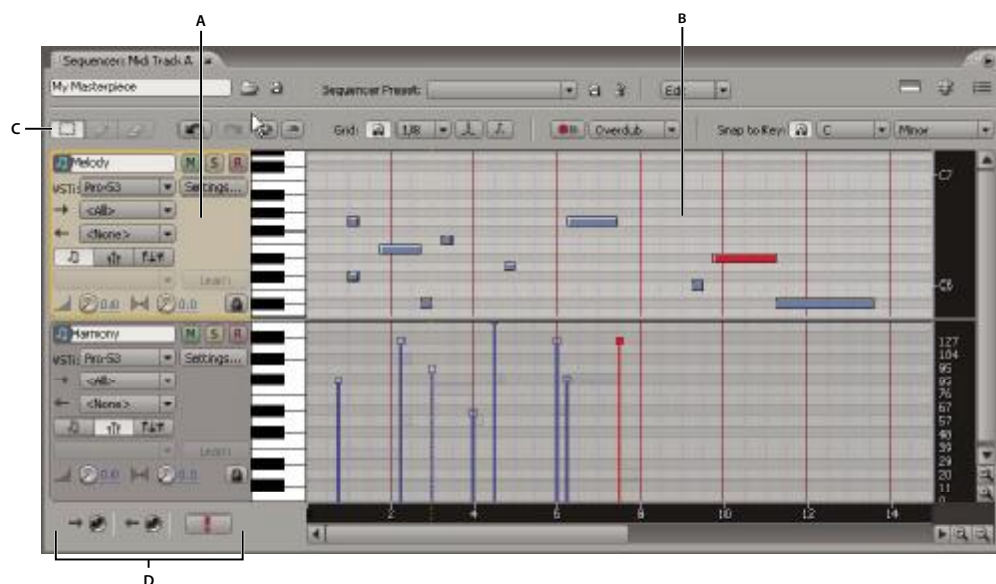
MIDI data includes information such as musical note, rhythm, and volume. When a MIDI file is played, this information is sent to a software or hardware instrument, which re-creates the notes.

Third-party VST instruments (VSTi's) provide hundreds of software-based synthesizers, samplers, drum machines, and more—without the expense and inconvenience of traditional hardware instruments. In the integrated mixing environment of Multitrack View, the MIDI tracks that host VSTi's behave much like audio tracks. You can apply effects and equalization, adjust volume and pan, and route the signal to sends and buses. Unlike audio tracks, however, you can also change the underlying musical data at any time.

Overview of the Sequencer

In the Main panel and Mixer, MIDI tracks contain a powerful sequencer that lets you manipulate MIDI data and interact with virtual and hardware instruments. Though each MIDI track contains its own sequencer, all synchronize to the same session properties. (See “Set tempo, time signature, and key for a session” on page 223.)

In the Sequencer, you'll find several unique options for recording and editing MIDI data on individual instrument tracks. Other options, such as Mute, Solo, Volume, and Pan, function like the similar options in Multitrack View.



The Sequencer

A. Track controls B. Note grid C. Editing tools D. MIDI activity lights and Panic button

See also

“Importing with the Files panel” on page 48

“Add, insert, or delete tracks” on page 182

“Mute and solo tracks” on page 184

“Create an audio clip from a bus, master, MIDI, or ReWire track” on page 201



Insert MIDI into a session and access the Sequencer

1 Do either of the following:

- Choose Insert > MIDI Track to insert an empty track.
- Choose Insert > MIDI to insert a MIDI file. Adobe Audition automatically creates a MIDI track and imports the file into the Sequencer.

2 In the track controls of the Main panel or Mixer, click Sequencer.



To add or remove instrument tracks in the Sequencer, click the Add Track  or Delete Track  button.

See also

“Export and import MIDI files in the Sequencer” on page 216

Enable VSTi plug-ins

The Sequencer can access any VSTi plug-ins installed in specified VST folders. (See “Enable VST effects” on page 110.) Within those folders, you must enable the individual plug-ins you want to load. Note that loading large numbers of plug-ins can cause Adobe Audition to start slowly.

1 In the Sequencer, click the VSTi Plug-In Manager button .

2 In the Enable column, select each plug-in you want to load. Or click Enable All or Disable All.



3 Click Rescan, and then click OK.

Assign, adjust, and mix VST instruments


After you enable VSTi plug-ins, you can assign virtual instruments to sequencer tracks. To tailor sounds for your compositions, customize instrument options.

1 In the track controls, select an instrument from the VSTi pop-up menu.

2 To display the instrument interface, click Settings, and adjust options as desired. (For information about options, see the documentation provided by the instrument developer.)

3 In the track controls, adjust Volume  and Pan  knobs to fine-tune the mix of VSTi's in the Sequencer. (To adjust the entire MIDI track, use controls in the Main panel or Mixer.)





MIDI tracks share monitoring options with audio tracks. (See “Record audio clips in Multitrack View” on page 57.) To monitor VSTi's, you must click the Live Monitoring button  of the related MIDI track in the Main panel or Mixer.

Use sequencer presets

Sequencer presets store settings for multiple tracks, providing a quick way to load an entire rack of instruments, assign MIDI inputs and outputs, and set volume and pan. A set of default presets gets you started; add your own to include third-party VSTi's or automatically connect to your MIDI hardware.

❖ Do any of the following:

- To load a preset, select it from the Sequencer Preset menu.
- To delete the currently selected preset, click the Delete Preset button .
- To save current settings for all tracks as a preset, click the Add Preset button .

See also

“Use effect presets” on page 104


Importing, recording, and outputting MIDI

Connect to MIDI inputs and outputs

MIDI inputs connect to devices like keyboards and control surfaces, so you can record MIDI data. MIDI outputs connect to devices like synthesizers and samplers, so you can play MIDI data.



Because the Sequencer supports software-based VST instruments, a MIDI output device isn't necessary for playback.

- 1 In the Sequencer, click the MIDI Devices button .
- 2 Select the inputs and outputs you want to access from sequencer tracks, and click OK.

Record MIDI in the Sequencer



The Sequencer lets you record MIDI from both external controllers and a built-in, virtual keyboard. If you can confidently play a musical passage, you can record a continuous MIDI performance, much as you would an acoustic performance. If a musical passage is particularly complex, however, you can record in steps defined by the current Grid settings.

Record from an external MIDI device

If you have a hardware controller like a keyboard, you can use it to record MIDI data in real time, while hearing an entire session.



- 1 Connect the device to a MIDI input.
- 2 In the track controls, select a MIDI channel from the Input menu.

The selected MIDI channel receives data from all inputs selected in the MIDI Devices dialog box. (See “Connect to MIDI inputs and outputs” on page 215.)




- 3 Click the Arm For Record button  for the sequencer track.
- 4 Click the Record button  in the Transport panel.
- 5 Play the MIDI controller.

Record with the virtual keyboard

If a hardware controller is unavailable, you can monitor VSTi's and record MIDI data with the virtual keyboard.

- 1 In the upper-right corner of the Sequencer, click the Virtual Keyboard button.
- 2 Click the Arm For Record button  for the sequencer track.
- 3 Click the Record button  in the Transport panel.
- 4 Play the virtual keyboard.

Record in steps

- 1 To determine the timing of each step, choose a Note Value setting in the Grid options. (See “Determine musical timing with the MIDI grid” on page 217.)
- 2 Click the Step Record button .
- 3 Click the Arm For Record button  for the sequencer track.
- 4 Click the Record button  in the Transport panel.
- 5 Play either a hardware controller or the virtual keyboard.

Set the record mode

Choose one of the following from the Record Mode pop-up menu:

Overdub Adds to existing data on tracks.

Replace Supplants existing data on tracks.

Output sequencer tracks to MIDI hardware

If you have hardware sound modules, you can output MIDI data directly to them from the Sequencer.

- ❖ In the track controls, choose a MIDI channel from the Output menu.



The selected MIDI channel transmits data to all outputs selected in the MIDI Devices dialog box. (See “Connect to MIDI inputs and outputs” on page 215.)



To include the audio output of a sound module in exported mixdown files, record the output on an audio track.

Export and import MIDI files in the Sequencer

To quickly share sequencing data with other multitrack sessions or sequencing applications, export and import MIDI files. These files contain all note, velocity, and controller information. After importing them, you simply need to assign a VSTi or hardware output to each track.

- ❖ In the Sequencer, do either of the following:
 - To export MIDI, click the Export MIDI File button .
 - To import MIDI, click the Open MIDI File button .



Note: Imported MIDI replaces any existing MIDI in the Sequencer.


See also

“Insert MIDI into a session and access the Sequencer” on page 214

Monitor or stop all MIDI activity

In the lower-left of the Sequencer, do either of the following:

- To confirm that a MIDI port is receiving or sending data, look for a red light in the MIDI In Activity or MIDI Out Activity indicator .
- If you hear stuck notes or other unexpected MIDI behavior, click the Panic button  to instantly stop playback of all MIDI data.

 *If you experience odd MIDI behavior outside of the Sequencer, choose Options > MIDI Panic Button.*

See also

“Connect to MIDI inputs and outputs” on page 215

Displaying and editing MIDI

Zoom and scroll MIDI

In the Sequencer, vertical zoom options are independent from those in the Main panel, letting you display musical notes and audio amplitude with differing levels of detail. By contrast, horizontal zoom options link the two panels, ensuring that both display the same time range so you can tightly synchronize edits.


- To vertically zoom a track, position the mouse over the vertical keyboard or data ruler, and roll the mouse wheel.
- To vertically scroll a track, drag the vertical data ruler at the right of the track up or down.
- To vertically resize a track, drag the track borders up or down.
- To fit all note data in a track, double-click the MIDI grid.
- To vertically zoom all tracks, click the zoom buttons below the vertical scroll bar.
- To horizontally zoom all tracks, click the zoom buttons to the right of the horizontal scroll bar.


See also



“Viewing, zooming, and navigating audio” on page 19

Determine musical timing with the MIDI grid

Grid options let you precisely evaluate and adjust musical timing. These options determine the default length of manually entered notes, as well as intervals for step recording. Grid snapping ensures that edited notes synchronize with the beat.

- ❖ In the Grid options area of the Sequencer, do any of the following:
 - Click the Snap To Grid button .
 - From the Note Value menu, choose a beat subdivision (for example, 1/4 or 1/8).


 *The Sessions Properties panel determines the number of beats per bar. (See “Set tempo, time signature, and key for a session” on page 223.)*

- Click the Triplet button  to divide each beat into three subdivisions. This option is particularly useful for music with swing and shuffle rhythms, like jazz and blues.
- Click the Dotted button  to mark every 1.5 beats. This option is particularly helpful for music with an odd number of beats per bar. In a waltz with three beats per bar, for example, a dotted note perfectly spans half a measure.

Snap edited or recorded MIDI to musical scales

In the Sequencer, Snap To Scale options restrict all newly edited and recorded notes to a musical scale you specify. Snapping to a scale gives music a unified, harmonic feel.

By default, the root note of the scale reflects the session key, which controls the pitch of all voiced audio loops in a session. If your session doesn't contain loops, or if you understand how musical keys interact, feel free to specify a root note that differs from the session key.

- 1 Click the Snap To Scale button .
- 2 From the pop-up menus to the right of the button, select a root note and scale type.

See also

“Set tempo, time signature, and key for a session” on page 223

Edit MIDI notes




The Sequencer lets you edit notes with an intuitive, grid-based approach that even nonmusicians easily understand.

See also


“Edit velocity and controller data” on page 219

“Prevent changes to MIDI data” on page 220



Select MIDI data

The Sequencer provides a variety of ways to select MIDI data on the grid. Selections in a track's Note  and Velocity  views are automatically linked. Selections in the Controller view  are independent.


Do any of the following:

- Click or drag with the Select tool . (Hold down Ctrl to select additional notes.)
- Click the desired track, and then choose Edit > Select All. (To add notes on another track, click that track, and again choose Edit > Select All.)
- In the vertical keyboard for a track, click a key to select all notes at that pitch.

Delete notes




❖ Click individual notes with the Erase tool . Or, select groups of notes with the Select tool , and then choose Edit > Delete.

Manually enter notes




❖ Click the grid with the Draw tool .

The initial length of newly-entered notes is defined by the Note Value setting in the Grid options. (See “Determine musical timing with the MIDI grid” on page 217.)

Trim or extend notes

- 1 Click the Select tool  or the Draw tool .
- 2 In the grid, position the pointer over the left or right edge of a note. When the pointer becomes a double-arrow , drag.

Move notes

- 1 Click the Select tool  or the Draw tool .
- 2 In the grid, position the pointer over the middle of a note. When the pointer becomes a hand , drag.

Copy and paste notes

- 1 Select the notes in the grid. Then choose Edit > Copy.
- 2 Click the destination track at the desired start point, automatically moving the current-time indicator. Then choose Edit > Paste.

Pasted notes initially reflect original pitch values, but you can move pasted notes if needed.



To quickly copy selected notes to another location in the same track, Ctrl-drag them.

Automatically adjust timing, velocity, and transposition

- 1 In the grid, select the desired data.
- 2 From the Edit menu, choose any of the following:
 - Humanize** Slightly offsets timing relative to the grid, providing a human feel to notes you’ve manually entered.
 - Quantize** Aligns note timing with the grid, a useful option for recorded performances that are rhythmically out of sync.
 - Randomize Velocity** Assigns arbitrary velocity values to each selected note, often producing a more human feel, particularly for drum tracks.
 - Transpose** Adjusts pitch up or down in semitones (half-steps) and full octaves that you specify in a separate window.

Undo or redo edits

- ❖ In the upper-left of the Sequencer, click the Undo  or Redo  button.

Edit velocity and controller data

MIDI velocity data determines how hard each note is struck; MIDI controller data adjusts common synthesizer parameters such as modulation, filter resonance, and portamento. To create controller data, you adjust assigned knobs and sliders on a MIDI hardware controller while recording performances. You can then edit the recorded values in the Sequencer, fine-tuning them to achieve precise musical results.


Both velocity and controller data have value scales ranging from zero to 127. At a velocity of zero, for example, a note is unplayed and inaudible. At the maximum velocity of 127, however, a note is played as loudly as possible.

See also


“Edit MIDI notes” on page 218

“Prevent changes to MIDI data” on page 220


Edit velocity data

- 1 In the track controls, click the Velocity button .
- 2 In the grid, drag the top of velocity lines up or down. (Ctrl-click to select multiple lines and adjust them simultaneously.)

Assign a hardware control to a parameter


- 1 In the track controls, click the Controller button .
- 2 From the pop-up menu, select the parameter.
- 3 Click Learn.
- 4 On the MIDI hardware controller, move the knob, slider, or wheel you want to control the parameter.

Edit controller data

- 1 In the track controls, click the Controller button .
- 2 From the pop-up menu, select the parameter.
- 3 In the grid, drag the top of controller bars up or down. (Ctrl-click to select multiple bars and adjust them simultaneously.)

Prevent changes to MIDI data

When editing a particularly complex MIDI composition, you might inadvertently change existing note, velocity, or controller data on sequencer tracks. To protect a track from changes, lock it.

- ❖ In the track controls, click the Lock button .

See also

“Freeze effects and improve performance” on page 109

Chapter 11: Loops

Using loops in Multitrack View


About loops

You can hear loop-based compositions in nearly all musical styles. From best-selling pop, rap, and hip hop, to alternative, adult-contemporary, and jazz, you'll find loops used for everything from basic rhythm tracks to entire compositions. With Adobe Audition, you can either create your own loops or choose from thousands supplied on the Adobe Audition Loopology DVD.

Loops let you create extremely flexible multitrack sessions. You can incorporate a single loop into many different compositions, each with adjustable musical key and tempo. Though loops typically contain only one or two bars of music (four to eight beats), you can extend and repeat them by simply dragging with the mouse. With snapping enabled, you can even use this method to rhythmically synchronize loops. For example, you can drag to create 1.5 repetitions and end precisely on a snare hit at a loop's midpoint.



For tutorials, training, and other resources, visit Resource Center on the Adobe website.



Identifying loops in the track display
A. No loop B. Single loop C. Extended (repeated) loop

Rhythmically synchronize loops

To rhythmically synchronize a loop-based session, use the Bars And Beats time format and enable snapping. Then add loops to create a rhythmic foundation, which you can build upon by recording new audio clips. (You can also add existing audio clips, but only their start and end points will align with loop beats.)

1 Choose View > Display Time Format > Bars And Beats. (This ruler format makes it easier to visually align loops with musical beats.)

2 From the Edit > Snapping submenu, choose any of the following:

Snap To Ruler (Coarse) To snap to beats within bars. Use this option if you work with 1/4 or 1/2 bar loop files.

Snap To Clips To snap to the start or end of audio clips.

Snap To Loop Endpoints To snap to the start or end of loops.


See also

“Snap to markers, rulers, frames, and zero crossings” on page 76

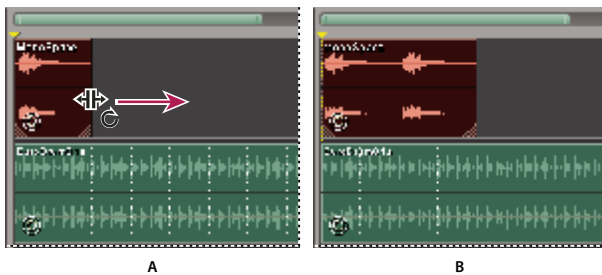
“Snap to clip and loop endpoints” on page 193

“To change the time display format” on page 55

Extend or shorten a looped clip

- 1 Select the clip, and then position the pointer over the bottom left or right handle—the loop editing  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. Positioning the pointer over the clip handle B. Dragging the loop, with snap-to lines indicating beats in other tracks

Override original loop properties

Loop properties that you set in Multitrack View are saved with sessions, overriding original loop properties without changing source files. (See “Set original loop properties” on page 225.)

Note: By default, property changes in Multitrack View affect only the selected clip, unless you select *Adjust All Loop-Enabled Clips That Use This Wave*.

- 1 In Multitrack View, select an audio clip.
- 2 Choose Clip > Loop Properties.
- 3 In the Audio Clip Looping dialog box, set the following options, and click OK:

Enable Looping Lets you extend a looped 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 the 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. If you deselect 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 beat. If you change tempo, the audio clip moves so that it starts at the same beat. Select this option if you're stretching a loop to tempo or starting a one-shot clip (like a thunder clap or a gong) on musical beats.

Source Waveform Information Lets you override original tempo and key properties stored in the source file.

Tempo Matching Lets you override original stretch method properties stored in the source file.

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 source file. For example, if you insert the same loop file into a session four times and then adjust the loop properties of one of the four clips, the other three instances are also adjusted.

Set tempo, time signature, and key for a session

The Session Properties panel lets you specify the tempo, time signature, and key for audio loops in a session. All loop-enabled clips automatically adjust to match new settings; regular clips are unaffected.



Session tempo and time signature also affect MIDI tracks. (See "Overview of the Sequencer" on page 213.)

1 In Multitrack View, choose Window > Session Properties if the panel 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 musical key for voiced audio loops.

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 you can customize properties such as mixing and metronome settings.

Metronome Toggles the built-in metronome on and off.



To preview loop files at the tempo and key of a session, select the Loop option in the Insert Audio dialog box, or the Follow Session option in the Files panel. (See "Insert an audio file into a session" on page 47 and "Preview audio in the Files panel" on page 49.)

See also

"Calculate the tempo of a selected range" on page 225

"Set advanced session properties" on page 210

Creating loops in Edit View

Select a range for a loop


To create a good loop from an existing file, you first need to select a range that repeats precisely on a beat. Repeating on a beat lets you combine a loop in rhythm with other loops.

- 1 In Edit View, open a file from which you want to create a loop.
- 2 Choose Edit > Auto-Mark > Find Beats And Mark. Then specify the Decibel Rise and Rise Time for beats, and click OK.
- 3 Choose Edit > Snapping > Snap To Zero Crossings.


This command makes your selection snap to places in the waveform that have zero amplitude, preventing audible noise at the beginning and end of the loop.

- 4 Select the range that you want to become a loop, typically starting and ending on a beat.

For guidance, use the beats you marked in step 2, together with commands in the Edit > Find Beats menu. (See “Select audio between beats” on page 75.)



Selecting a range that starts and ends on a clear beat helps create a good loop.

- 5 Click the Play Looped button  to repeatedly play your selection.
- 6 Adjust the start and end of the selection to fine-tune the loop.
- 7 Choose View > Display Time Format > Edit Tempo to calculate the tempo of the selected range, and then click OK.
- 8 Choose Edit > Copy To New.
- 9 In the new file, choose File > File Info, click the Loop Info tab, and set original loop properties.



To create a loop that repeats seamlessly, see “View the top and tail of an audio file” on page 71.

See also

“Snap to markers, rulers, frames, and zero crossings” on page 76

“Calculate the tempo of a selected range” on page 225

“Set original loop properties” on page 225

Calculate the tempo of a selected range

To calculate the tempo of a selected range, use 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.

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 extract 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, a time setting of 4:2:16 represents an eighth note halfway between beats 2 and 3 in 4/4 time.

Set original loop properties

For either new or existing loop files, you can set loop properties such as tempo and key. Adobe Audition uses these original properties to accurately adjust loops for different session tempos and keys.

1 In Edit View, choose File > File Info.

2 Click the Loop Info tab.

3 Set any of the following options, click OK, and then save the file:



For the best audio quality, save loops in uncompressed formats, such as Windows PCM (.wav).

Loop Automatically enables looping if the file is inserted into a multitrack session.

One Shot Indicates that the file plays once rather than repeating 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 original key so Adobe Audition can adjust it to match each session's key. For a drum track, choose Non-Voiced—unless you want to pitch-shift drums when you change session key.

Find Nearest Scans the loop to locate the nearest key. This option works best with 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 options:

- **Fixed Length (No Stretching)** retains clip length without time-stretching to match either session tempo or key. Instead, this option resamples—changing tempo and key simultaneously. Use this option only if you want to duplicate a clip by dragging. (Note, however, that the Edit > Clip Duplicate command provides similar functionality without unexpected tempo and key changes.)
- **Time-Scale Stretch** (which reflects the Stretch effect) extends the loop to match the tempo of the session. Options include Quality, Solo Voice/Instrument, and Preserve Formant. (The latter two options maintain realism during pitch shifts.) This method stretches a file based on its actual length, so you should use it to loop material like synth pads or sustained string sections, which don't have well-defined beats. 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 if you're trying to change the sound and timbre of vocals.
- **Beat Splice** loops the file based on beats detected in it, similarly to the Find Beats And Mark command. (See “Select a range for a loop” on page 224.) 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.

Chapter 12: Working with video

Working with video applications

Edit audio clips from After Effects

In Adobe After Effects®, choose Edit > Edit In Adobe Audition to open audio clips in Edit View, where you can quickly restore and enhance them. When you save your changes, the updated audio automatically appears in your video projects.

For more information, search for “Adobe Audition” in After Effects Help.

See also

“Techniques for restoring audio” on page 141

“Mastering effect” on page 161

Link sessions to exported WAV mixdowns

To create layered video soundtracks that you can easily update, link multitrack sessions to exported WAV mixdown files. When you select the exported files in an application such as Adobe Premiere Pro, you can either remix or edit them in Adobe Audition. As a video project evolves, simply repeat this process to create a polished final soundtrack.

- 1 Open a multitrack session.
- 2 Choose File > Export > Audio Mix Down.
- 3 In the Mix Down Options section, select Embed Edit Original Link Data.
- 4 Specify a filename, choose a WAV format, and click Save.

See also

“Comparing Edit View and Multitrack View” on page 20

Edit a WAV mixdown via a video application

- 1 In the video application, select the WAV mixdown you exported from Adobe Audition.
- 2 Choose Edit > Edit Original.
- 3 Select one of the following options, 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 To overwrite the original file in the video application, do 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.

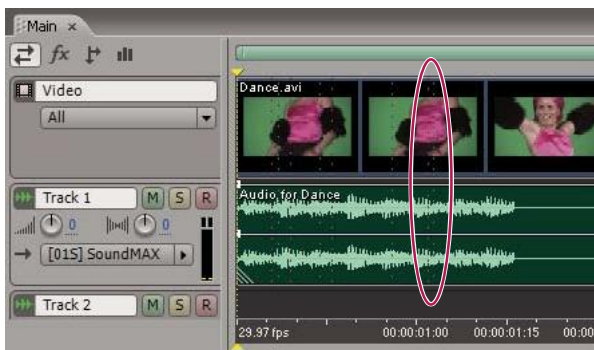
Importing video and working with video clips

Insert a video file into a session

In Multitrack View, you can insert video files to precisely synchronize a session with a video preview. When you insert a video file, its video clip appears at the top of the track display, and its audio clip appears on the track below. For AVI files exported from Adobe Premiere Pro, sequence markers appear as dashed lines.


You can move a video clip independently from the audio clip containing the original soundtrack. To instead keep such clips synchronized, group them. (See “Grouping clips” on page 192.)

Note: A session can contain only one video clip at a time.



Related video and audio clips, with imported video markers

- 1 In Multitrack View, position the start-time indicator at the desired insertion point.
- 2 Choose Insert > Video, select an AVI, QuickTime, or Windows Media file, and click Open.

 To import only the audio from a video file, choose Insert > Audio From Video (Multitrack View) or Open > Audio From Video (Edit View). This technique is great for soundtrack editing that doesn't require a video preview, or for readapting soundtracks for audio-only mediums, such as radio or CD.

See also


“Video file formats” on page 252

“Select and move clips” on page 192

Snap to frames in a video clip

To synchronize a multitrack session with video, snap audio clips and the start-time indicator to frames.

- 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 zoom the session display. Be aware, however, that video thumbnails serve only as a general guide; for frame-accurate synchronization, use snapping.



Snapping to a video frame that falls within a video thumbnail

See also

“Snap to markers, rulers, frames, and zero crossings” on page 76

“Snap to clip and loop endpoints” on page 193

“Zooming audio” on page 20

Change how a video clip displays thumbnails

You can change how a video clip displays thumbnails of frames. Displaying more thumbnails helps you identify video content; displaying fewer thumbnails increases performance.

❖ In the video track controls, choose one of the following options from the Thumbnail Options menu:

No Thumbnail Displays a blank video clip.

First Only Displays a thumbnail for only the initial frame.

All Displays thumbnails for every frame. (The session’s horizontal zoom setting determines the percentage of visible frames.)

See also

“Zooming audio” on page 20

Previewing video

Video panel overview

In the Video panel, 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 Video panel to optimize it for your monitor size and system speed. For example, you can scale previews to fit a resized panel or lower preview quality to increase performance.



To hide or show the Video panel, choose Window > Video.



Video panel and related clip in Main Panel

To customize the Video panel, right-click it, and select any of the following options:

A zoom percentage Zooms in or out.

Best Fit Fits previews to the panel.

Maintain Aspect Ratio Maintains aspect ratio when you resize the panel.

Integer Factor Sizing Constrains previews to ratios such as 1/2, 1/1, and 2/1 when you resize the panel. This option avoids complex resampling, producing a sharper image and increasing performance.

Auto Show Video Automatically shows the Video panel when you insert video files.

Low Quality Lowers the preview quality, increasing performance.

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.

Preview exported video


In Multitrack View, you can export a video file in the format of an imported video clip. If you import an AVI file, for example, you can export only in AVI format.

Exported video combines a video clip with audio clips in the same selected range, creating a new soundtrack. Prior to exporting video, preview it to ensure that it will sound as you expect, and then edit the session if needed.



Selecting start and end points for exported video

- 1 Choose Edit > Snapping > Snap To Frames.
- 2 In the session display, select a range that extends from the beginning to the end of the video clip.
- 3 Play the session, and then do one of the following:
 - If you like the soundtrack, export a video file (choose File > Export > Video).
 - If you want to adjust the soundtrack, edit the session as needed, and then repeat steps 2 and 3. (For example, if part of an audio clip is omitted, move the entire clip into the selected range.)

 You can also use this procedure to export an audio mixdown that you recombine with video in a video application, such as Adobe Premiere Pro. Though exported video files are limited to stereo audio, exported audio mixdowns support stereo and surround sound.

See also

- “Insert a video file into a session” on page 228
- “Export a session to a video file” on page 243
- “Export a session to an audio file” on page 242
- “About surround sound” on page 232

Chapter 13: Creating surround sound

Surround sound basics

About surround sound

With surround sound, heard in many popular movies, you can pan an audio mix around a 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.


To accurately monitor low frequencies, configure the LFE subwoofer to reflect your audience's surround system. For example, Digital Theater System (DTS) typically uses a crossover frequency of 80 Hz, routing all content below that frequency to the subwoofer. However, most Dolby® Digital systems use a crossover of 120 Hz. Some systems also boost the LFE channel by 10 dB, emphasizing low-frequency content.

Create mixes with the Surround Encoder

You export 5.1 surround sound from a multitrack session by using the Surround Encoder dialog box. With this dialog box, you can pan each track, preview the surround mix, and export it as six mono WAV files, or one interleaved, six-channel WAV or WMA file.

Important: To accurately preview surround mixes, you need a sound card with at least six outputs. For detailed requirements, see “Set the preview device for surround mixes” on page 237.

1 In Multitrack View, open or create a session, and create a basic stereo mix with desired track volume, pan, and effects settings.

 If you open an existing session, consider choosing *File > Save Session As* to save a copy specifically for surround mixing.

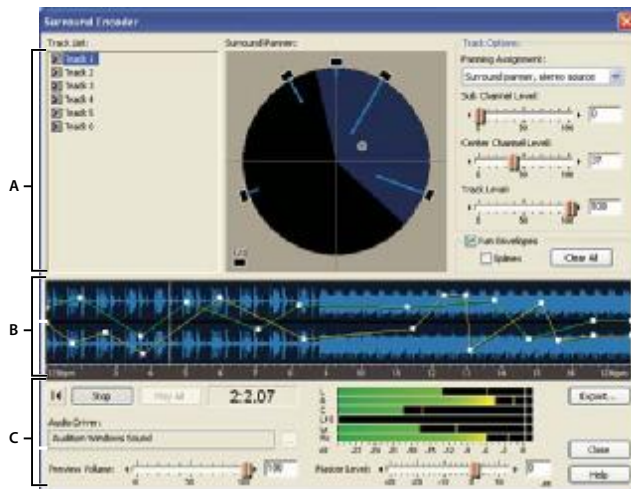
2 Choose *View > Surround Encoder*.

3 In the Track List area, select the tracks and buses you want to include in the surround mix.

4 In the Track Options area, set the panning assignment and amplitude levels for each track.

5 If necessary, fine-tune and automate panning for each track.

6 Preview and export the surround mix.



Surround Encoder dialog box
A. Track controls B. Waveform display with pan envelopes C. Preview and export controls

See also

“Include a track or bus in a surround mix” on page 233

“Assign a surround panning control” on page 234

“Adjust volume levels for surround mixes” on page 236

“Preview controls for surround mixes” on page 236


“Export a surround mix” on page 238

Setting track pan and volume for surround mixes

Include a track or bus in a surround mix

❖ In the Track List area of the Surround Encoder, click the check box next to the track or bus name.

💡 To access a track that outputs to a bus, send the track to an additional, track-specific bus. (See “Routing audio to buses, sends, and the Master track” on page 186.)



Including a track in a surround mix

Assign a surround panning control

❖ At the top right of the Surround Encoder dialog box, choose one of the following Panning Assignment options for each track:

Surround Panner, Stereo Source Retains stereo imagery by using the Surround Panner as a balance control, attenuating content in the left and right channels without moving it. For example, if the source track is stereo, the left channel outputs to the Front Left and Left Surround channels, and the right channel outputs to the Front Right and Right Surround channels. The Center channel receives a summed (L + R) mono signal.

Surround Panner, Summed To Mono Lets you use the Surround Panner to position signals anywhere in the surround field, but sums stereo tracks to mono.

LFE Only Sends the entire track to the LFE (subwoofer) channel. Your monitoring system should apply a crossover frequency that reflects the surround system your audience will use. (See “About surround sound” on page 232.)

Note: The Surround Encoder does not filter audio sent to the LFE channel when previewing, exporting, or encoding. Apply any necessary low-pass filtering to either source tracks in Multitrack View or exported files.

L + R, Stereo Sends the selected track’s signal in stereo to only the Front Left and Front Right speakers.

Ls + Rs, Stereo Sends the selected track’s signal in stereo to only the Left Surround and Right Surround speakers.

Center + LFE, Stereo Most useful for stereo tracks, this option routes the left signal to the Center channel and the right signal to the LFE (subwoofer) channel. (For mono tracks, the same signal is sent to both the Center and LFE channels.)

Center Only, Mono; L Only, Mono; R Only, Mono; Ls Only, Mono; Rs Only, Mono Sums the selected track’s audio to mono, and sends it all to the selected channel. (You can also access these options by dragging the pan point onto one of the five main speakers in the Surround Panner.)

See also

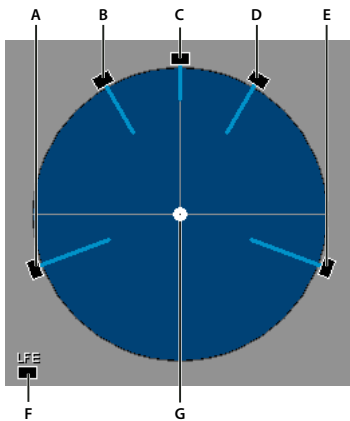
“Use the Surround Panner” on page 234

Use the Surround Panner

❖ In the Surround Encoder dialog box, drag the pan point (white dot) to change the position of the signal.

As you move the pan point, the light blue lines coming from the speakers change length. These lines reflect the power of the signal in each of the five main channels. In the background, a dark blue area of the sphere rotates, expands, and contracts. This area reflects the image of the surround field for a listener seated in the center of the speakers (marked by the crosshairs).


You can also drag the pan point directly over any 5.1 speaker, summing audio from the selected track to mono, and sending it to only one speaker channel.



Surround Panner options
A. Left Surround B. Front Left C. Center D. Front Right E. Right Surround F. Low Frequency Effects (Subwoofer) G. Pan point

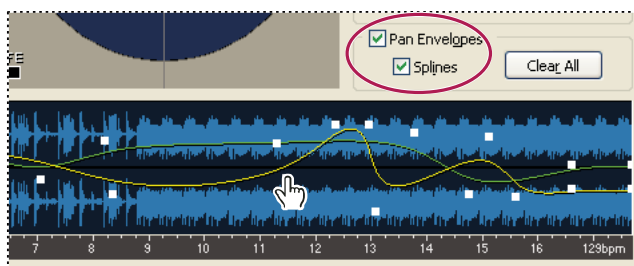
Statically pan tracks in a surround mix

- 1 In the Surround Encoder dialog box, select a track, and deselect Pan Envelopes.
- 2 In the Surround Panner, set the pan point to any static position.

 If you toggle a track's Pan Envelopes option on and off, any envelope points you have created are retained.

Dynamically pan tracks in a surround mix

When you select Pan Envelopes in the Surround Encoder, 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 interact with the pan point in the Surround Panner and let you dynamically pan tracks over time.



Pan envelopes in the Surround Encoder

- 1 In the Surround Encoder dialog box, select 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 Above the right side of the waveform display, select Pan Envelopes.

Two envelope lines appear in the waveform display. (Because the yellow line starts on top, you won't see the second, green line until you change the pan position.)

- 4 Do any of the following:
 - To add a control point in the waveform display, click an envelope line. Alternatively, click where you want to set a control point, without clicking an envelope, and then drag the pan point in the Surround Panner.

- To edit a control point, drag it. The pan point moves in tandem to show you the relative position in the surround field.
 - To delete an individual control point, drag it up or down beyond the waveform display.
 - To delete all control points, click Clear All.
 - To create smoother transitions between control points, select Splines.
- 5 Click Play Track or Play All to hear the dynamically panned track.

Adjust volume levels for surround mixes

❖ In the Surround Encoder dialog box, select a track, and then set any of the following options:

Sub Channel Level Controls the amplitude level sent to the LFE channel.



Use a bass management circuit in your monitoring setup to ensure that you hear the same bass levels as your audience.

Center Channel Level For tracks assigned to the Surround Panner, determines the Center channel level relative to the Front Left and Front Right level. At 100, the Center channel receives an equal percentage of signal.

Track Level Controls the amplitude level of the selected track.

Zoom waveforms in the Surround Encoder


❖ In the Surround 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 from the context menu.
- Right-click and drag the desired zoom area on the time ruler. (To zoom back out, right-click and choose Zoom Out or Zoom Full from the context menu.)
- Position the mouse pointer in the waveform display, and spin the mouse wheel. This zooms into the time area directly beneath the mouse pointer.

Previewing and exporting surround mixes


Preview controls for surround mixes

The preview controls in the Surround Encoder let you fine-tune a surround mix before exporting it.



Preview controls in the Surround Encoder

A. Transport controls B. Time indicator C. Output meters D. Preview device E. Preview volume F. Preview device change button G. Master Level

Go To Beginning  Places the cursor at the start of the track.

Play Track Plays the selected track from the cursor location to the end of the track, regardless of the current zoom level.

Play All Plays the complete surround mix from the cursor location to the end of the session, regardless of the current zoom level.

Time indicator Located next to the Play All button, this indicator shows the time at the current cursor location.

Preview Volume Controls preview volume without affecting the amplitude of exported files. (Use the Master Level slider to control exported amplitude.)

Output meters (L, R, C, LFE, Ls, Rs) This set of six meters displays the amplitude of each surround channel. If you click Play Track, the meters display the output of only the selected track. If you click Play All, the meters display the output of the complete 5.1 mix, reflecting levels for exported files.

Master Level Sets the overall amplitude for exported files. Use this slider with the output meters to maximize overall amplitude and ensure that no channels are clipping.

Preview Device, Format Displays the currently selected device to which Adobe Audition routes the six-channel, surround output.


See also

“Set the preview device for surround mixes” on page 237

Set the preview device for surround mixes

To accurately preview 5.1 surround mixes, you need a DirectSound or ASIO sound card with at least six outputs. If you use a DirectSound card, you need a special interleaved driver.

In the Surround Encoder, the Preview Device, Format box shows the current output driver and bit depth for surround mixes. You can change these settings in the Audio Hardware Setup dialog box.

- 1 Click the Audio Hardware Setup button  to the right of the Preview Device, Format box.
- 2 In the Surround Encoder tab, select an audio driver for the sound card you want to use, and then map each 5.1 channel to a sound card output. For more information, see “Set audio inputs and outputs” on page 31.

Note: Unlike ASIO sound cards, DirectSound cards that support 5.1 playback require a specific configuration of speakers, so you can't map channels to different outputs. However, some DirectSound cards provide different drivers for stereo and 5.1 output; in this case, select the 5.1 driver. (Cards like the Creative Labs Audigy, which uses one driver for both stereo and surround, will automatically route a 5.1 mix to the correct outputs.)

See also

“About surround sound” on page 232

“Create mixes with the Surround Encoder” on page 232


Export a surround mix

Adobe Audition can export surround mixes in either one WMA or WAV file that contains all six channels, or six separate WAV files for each channel.

Note: If you plan to deliver exported files to an outside service for surround encoding, confirm format requirements with that service.

- 1 Click Export at the bottom right of the Surround Encoder dialog box.
- 2 In the Multichannel Session Name box, enter a filename prefix for the exported files.

At the bottom of the dialog box, the Filenames To Be Saved section displays the complete filenames.



Complete filenames appear at the bottom of the dialog box.

- 3 In the Save In box, specify the folder for the saved files.
- 4 From the Format options menu, select one of the following:

Export As Six Individual Mono Wave Files Creates mono Windows PCM (.wav) files that can be used by any Windows audio application.

Export As One Interleaved, 6-Channel Wave File Creates a single, six-channel file in Windows PCM (.wav) format. However, some Windows audio applications support only mono and stereo WAV files.

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 a multichannel WMA file 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. Set 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** Determine how the surround mix will sound on a stereo system. Specify attenuation levels to mix the Center, Surround, and LFE channels with the front stereo channels. The defaults are good settings for most files, but you can enter custom values between 0 and -144 dB.

- **Show Codec Formats That Most Closely Match The Session's Sample Rate** Limits the list of selectable WMA kbps options to the sample rate and bit depth of files in the current multitrack session.

See also

“Windows PCM (.wav, .bwf)” on page 251

“Windows Media Audio (.wma)” on page 250

Chapter 14: Saving and exporting

Saving and exporting files

Save audio files individually

In Edit View, you can save audio files in a variety of common formats. The format you choose depends on how you plan to use the file. Keep in mind that each format stores unique information that might be discarded if you save a file in a different format.

1 Do one of the following:

- To save changes in the current file, choose File > Save.
- To save changes under a different filename, choose File > Save As.
- To save a copy of the current file while leaving the original version open and active, choose File > Save Copy As.
- To save currently selected audio as a new file, choose File > Save Selection.
- To save all open files in their current formats, choose File > Save All.



The Save Selection command is useful for saving small segments of a larger file. For example, you can use this command to break up a long recording into smaller, more manageable tracks.

2 Choose a location for the file, type a file name, and choose a file format. (See “About audio file formats” on page 244.)

3 Depending on the format you choose, additional options might be available. To view format-specific options, click Options.

4 Select Save Extra Non-Audio Information to preserve markers and file information, such as loop and Broadcast Wave metadata.

***Note:** Deselect this option only when you will use another application to burn the file to CD. Some CD-recording applications misinterpret non-audio information (such as markers and metadata), producing an unpleasant burst of noise at the beginning of each track.*

5 Click Save.

See also

“Add audio file information” on page 253

“About markers” on page 80

Save a group of audio files to one format



In either Edit View or Multitrack View, you can quickly save a group of open audio files to one format with the Save All Audio As command. This command automatically creates filenames from file information such as Artist, Album, and Song, providing a particularly efficient way to save groups of extracted CD tracks.


1 Choose File > Save All Audio As.

2 Set the following options, and then click OK.


Destination Folder Specifies the folder files are saved to. To visually navigate to a folder, click Browse.

Filename Template Names each file using information drawn from the Text Fields tab of the File Info dialog box. (See “Add audio file information” on page 253.) If you extracted a file from CD, any track information acquired from a CD database is available. Each symbol corresponds to information such as Artist (%a), Album (%d), and Song (%s).

To add a new filename template, click the plus  icon. To delete the current template, click the minus icon .

 *When adding a new template, you can specify subfolders within the destination folder; simply insert the subfolder name followed by a backslash (\). For example, the entry **MyFiles\%s** saves all files in a MyFiles subfolder and names them after the song title.*

Output Format Specifies the file format. To change settings for a selected format, click Options.

 *If the source files have different sample rates, bit depths, or channels, click Options to specify unique format settings for each sample type.*

Audio Files list In the Proposed Filename column, shows how saved filenames will appear. In the Status column, indicates filename conflicts and if files will be newly saved, overwritten, or skipped. If the Status field indicates Invalid Filename or Duplicate Filename, you must choose a different filename template. The current one specifies either unavailable or identical file information.

Show Folder Displays the complete path for each file in the Audio Files list.

Skip Excludes selected files from the saving process.

Overwrite Existing Files Automatically replaces files of the same name, without asking you to confirm each replacement.

See also


“Extract CD tracks with the Extract Audio From CD command” on page 51

“Audio file formats” on page 244

Save multitrack sessions

A multitrack session file is a small, non-audio file. It merely stores information about locations of related audio files on your hard drive, the duration of each audio file within the session, the envelopes and effects applied to various tracks, and so forth. You can reopen a saved session file later to make further changes to the mix.

If you create multitrack compositions entirely in Adobe Audition, save session files in the native SES format. If you plan to share multitrack compositions with other applications, however, save sessions in XML format.

 *Third-party developers may provide support for additional formats by using the Adobe Audition Software Developer’s Kit (SDK). For more information, see www.adobe.com/go/learn_au_sdk.*

1 In Multitrack View, do one of the following:

- Choose File > Save Session to save changes to the current session file.
- Choose File > Save Session As to save changes to a different session file.
- Choose File > Save All to save the session file and all the audio files it uses.

2 In the dialog box, select a location for the file, and type a filename.

3 From the Save As Type menu, choose the session file format.

4 (Recommended) 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.

Important: To more easily organize, move, and archive all files in a session, select the option above.

5 (Optional) 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. If necessary, click Format Properties to access options for the selected format. (See “About audio file formats” on page 244.)

6 Click Save.

Convert the sample rate of a session

1 Choose File > Save Session As, choose a location for the file, and type a file name.

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, see “Convert the sample rate of a file” on page 98.


5 Click Save.

Export a session to an audio file

After you finish mixing a session, you can export all or part of it in a variety of common audio file formats. When you export, current volume, pan, and effects settings are reflected in the resulting file.



To quickly mix specific audio clips down to a single track, use the Edit > Bounce To New Track command. (See “Create a single audio clip from multiple clips” on page 201.)

1 If you want to export part of a session, use the Time Selection tool  to select the desired range.

2 Choose File > Export > Audio Mix Down.

3 In the Export Audio Mix Down dialog box, specify a location, name, and format for the saved file. (See “About audio file formats” on page 244.)

If the file format you select can be customized, the Options button is available; click it to review or change settings, and then click OK.

4 In the Mix Down Options section, set Source, Bit Depth, and Metadata options.

5 Click Save.

See also

“Create an audio clip from a bus, master, MIDI, or ReWire track” on page 201

Audio Mix Down options

The following options are available on the right side of the Export Audio Mix Down dialog box.

Source Specifies the source of the file you are exporting:

- **Master** Exports the audio routed to the Master track.
- **Track** Exports a single track of the current session, which you select from the drop-down list.

- **Bus** Exports the audio routed to a specific bus.
- **Output** Exports the mix routed to the selected hardware output. By default, the output for the stereo master is selected; additional outputs are available if you have configured them in the Audio Hardware Setup dialog box and routed tracks or buses to them. (See “Assign audio inputs and outputs to tracks” on page 186.)

Range Specifies whether to mix down the entire session or a selected range. To mix down a range, you must select it before you choose the Audio Mix Down command.

Bit Depth Specifies 32- or 16-bit depth. If you select 16-bit, Enable Dithering and Dither Options are available. (See “Change the bit depth of a file” on page 99.)

Channels Specifies whether the mix down is mono or stereo.

Embed Edit Original Link Data Stores the path to the original session file, effectively linking related session and mix down files for Adobe Premiere and Adobe After Effects users. (See “Link sessions to exported WAV mixdowns” on page 227.)

Include All Markers and Metadata Saves header fields containing file information and markers in the file. (See “Marker types” on page 81.) If you plan to use Adobe Audition to burn the file to CD, leave this option selected. If you plan to use a different application, deselect this option, because that application may misinterpret non-audio information (such as markers and metadata), producing an unpleasant burst of noise at the beginning of each track.

Insert Mixdown Into Inserts the saved mix down file in Adobe Audition:

- **Edit View** Opens the mix down file in Edit View.
- **Multitrack View Audio Track** Creates a new audio track immediately below the currently selected track, and inserts the mix down file at the start-time indicator.
- **CD View** Inserts the mix down file into the current CD list, or creates a new CD list if one does not already exist. If the session contains track marker ranges, the ranges are automatically inserted as separate CD tracks.

Export a session to a video file

If a session includes a video file, you can mix down the session to create a new soundtrack in the same format as the video file (such as AVI, MOV, or WMV).

- 1 Choose File > Export > Video.
- 2 Do one of the following:
 - Choose a saved preset from the Preset menu.
 - Choose an option from the Audio Codec menu. Then select Mono or Stereo, and choose a sample rate from the Frequency menu.
- 3 Click OK.

See also

“About video file formats” on page 252

“Insert a video file into a session” on page 228

“Preview exported video” on page 230

Close files

❖ Do any of the following:

- To close the current audio file in Edit View, choose File > Close.
- To close the current session file in Multitrack View but leave related media files open, choose File > Close Session.
- To close a CD list in CD View, choose File > Close CD List.
- To close all audio and video files not in use, choose File > Close Unused Media.
- To close all open audio, video, session, and CD list files, choose File > Close All.

Audio file formats

About audio file formats

Adobe Audition lets you open and save files in a wide variety of audio formats. In most cases, you should save uncompressed audio in Windows PCM format and compressed audio in either mp3PRO or Windows Media Audio format. 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 them.

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 245.

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.

Choose from the following options:

A-Law 8-bit A slight variation of the standard mu-Law format and is found in European systems.

mu-Law 8-bit 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.

Choose from the following options:

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 “Change the bit depth of a file” on page 99.

Apple AIFF (.aif, .snd)

AIFF is the Apple® standard audio 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 cross-platform compatibility in both Windows and Mac OS. 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 from or written to 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: 1582BITSPERSAMPLE: 16  
CHANNELS: 2  
SAMPLERATE: 22050  
NORMALIZED: FALSE  
164 <tab> -1372  
492 <tab> -876
```

Choose any of the following options:

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 247.

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.

Choose one of the following options:

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 (Adaptive Differential Pulse Code Modulation) format is used in telephony applications, and it's optimized for voice at low sample rates. 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.

Choose from the following options:

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.

Choose from the following options:


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: simple 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.

Simple mp3 options

CBR (Constant Bitrate) 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 Bitrate) 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 with PRO data that 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 mp3 options

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 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 Combines 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 the padding required by the decoder. 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 AIFC, 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.

Choose from the following options:

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).

Ogg Vorbis (.ogg)

The Ogg Vorbis format is comparable to other formats used to store and play digital music, such as MP3, VQF, and AAC. Unlike those formats, however, Ogg Vorbis is license-free, so it's often used for commercial video games.

When you save an OGG file, you can either select one of three basic encoding options, or use advanced settings for detailed control:

VBR (Target Bitrate) Lets you specify the target bitrate in kilobits per second. Maintains audio quality by varying the bitrate depending on the complexity of the audio being encoded. This method can maintain higher audio quality, although file size is not as predictable as with Fixed Bitrate encoding.

VBR (Quality Index) Like VBR (Target Bitrate), but lets you specify quality on a scale of 0 to 10.

Fixed Bitrate Varies the quality level as needed to ensure that the bit rate stays at the specified rate. This method makes a consistently sized file, although the quality may not be as high as with Variable Bit Rate encoding.

Use Advanced Settings Enables the following options:

- **Minimum, Target, and Maximum Bitrate** Let you precisely specify compression settings.
- **Bit Reservoir Size** Specifies the amount of surplus bits to reserve during variable bitrate encoding.
- **Bit Reservoir Bias** Determines how surplus bits are distributed. Lower settings store surplus bits during consistent audio, instead applying those bits to transient peaks and troughs. Higher settings store surplus bits during transients, instead applying those bits to consistent audio. The default setting, 0.2, slightly favors transients.
- **Impulse Noise Floor** Sets the amplitude above which the encoder looks for artifacts in transient peaks. Lower noise floors improve transient response but increase bitrate.
- **Damping Time** Determines how quickly bitrate returns to the targeted average. At lower settings, bitrate varies less, but audio quality suffers; at higher settings, bitrate varies more, but audio quality improves.
- **Lowpass Filter** Specifies the highest frequency to retain in the encoded file.

SampleVision (.smp)

The SampleVision format is native to the Turtle Beach 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 Marker List panel. The Label of the marker 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).

Spectral Bitmap Image (.bmp)

Though most applications store conventional images in BMP format, Adobe Audition can save and import audio in this format. When you save audio as a bitmap image, Adobe Audition creates two image files: one reflects the spectral frequency graph, the other stores data that correctly aligns phase when you reimport the graph. (The latter file includes *phase* in the filename.) You can incorporate exported graphs into visual presentations, or modify them in image-editing applications like Adobe Photoshop.



In an image-editing application, apply gradients to create audio fades, and adjust opacity to change audio amplitude. To protect your work, apply visual watermarks above the audible frequency range.

For information about importing spectral graphs, as well as photographs, logos, and other visually-oriented files, see “Import a bitmap image as audio” on page 44.

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, .bwf)

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.



You can include Broadcast Wave metadata in Windows PCM files. (See “Add audio file information” on page 253.)

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 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.)

4-byte PCM (type 1, 32-bit) Saves 32-bit audio as 32-bit integers, retaining 192 dB of dynamic range. However, signal-to-noise ratio in saved files is 144.5 dB, reflecting the 32-bit float format that Adobe Audition uses internally.

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 +/-8million), but the .wav BitsPerSample is set to 24 while BlockAlign is still set to 4 bytes per channel.

16.8 float – Obsolete/Compatibility 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 depth of 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.

Choose from the following options:

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.

Video file formats

About video file formats

Adobe Audition lets you import and export video files in AVI, MOV, or WMV format. When you export to video, you must save in the format of the video file you imported in the session. For example, if you import an AVI file, you can export to video only in AVI format.

See also

“Insert a video file into a session” on page 228

“Export a session to a video file” on page 243

Uncompressed AVI (.avi)

AVI (Audio Video Interleaved) is a multimedia file format for storing sound and moving pictures on Windows systems. AVI is a container format, meaning that it specifies how data is organized, but it is not itself a form of audio or video compression.

Uncompressed AVI is a good alternative to DV AVI or other compression schemes because there is no loss of data on exporting.

DV AVI (.avi)

DV AVI format includes the DV codec compression scheme. This Windows format is primarily used to exchange sound and video with a DV camera through a Firewire (IEEE 1394) port.

QuickTime (.mov)

The QuickTime file format is used to create, edit, publish, and view multimedia files. QuickTime supports many different types of compression. Like AVI, QuickTime AVI is a container format, meaning that it specifies how data is organized, but it is not itself a form of audio or video compression.

Although Apple Computer, Inc. developed this format, its use is not limited to Mac OS. However, Windows users must install QuickTime for Windows to work with QuickTime files.


Windows Media Video (.wmv)

Windows Media Video, Microsoft's proprietary compression format for motion video, is used for both streaming and downloading content via the Internet.

Adding file information

Add audio file information

The File Info command opens a dialog box in which you can see and add *metadata*—descriptive information that a computer can process and search—about the open audio file. For example, metadata can include history, loop, and Broadcast Wave information.

 For audio files in WAV, mp3, or CEL format, you can use the File > XMP Info command to add XMP metadata. This extensive metadata can be shared across databases, file formats, and platforms. (See “Add or edit XMP metadata” on page 255.)

- 1 In Edit View, choose File > File Info.
- 2 Click the tabs at the top of the dialog box to navigate between different sets of properties.
- 3 Set properties as desired, and click OK.

 To preserve file information when you save files, select *Save Extra Non-Audio Information* in the Save dialog box.

See also

“Set original loop properties” on page 225

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 ID3v2.3-compatible data.

Sampler options

The Sampler tab of the File Info 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 File Info 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.

Broadcast Wave options

The Broadcast Wave tab of the File Info dialog box lets you view and edit metadata for broadcast applications.

Important: To include Broadcast Wave metadata in a file, you must save in Windows PCM (.wav) format and select the Save Extra Non-Audio Information option. For more information, see "Save audio files individually" on page 240.

Include BWF Data When Saving To .Wav Specifies that you want to include Broadcast Wave metadata when you save the file in Windows PCM (.wav) format. You must select this option in order to edit any of the other options. If you deselect this option, Adobe Audition deletes all Broadcast Wave metadata associated with the file when you click OK to close the dialog box.

Description Describes the audio file in up to 256 characters.

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 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 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 Specifies the timecode of the audio file, calculated since 12:00 a.m. (midnight). Select from the hh:mm:ss.ddd or Samples options. This option is commonly referred to as the clip's timestamp. It is used by Audition for spot inserting in a multitrack session.

Coding History Provides a text box for you to describe all coding processes applied to the waveform. Adobe Audition automatically adds information every time BWF data is modified and the file is saved. You can manually edit this information.

Write All Metadata Specifies whether to write metadata At the Start of the File or At the End of the File. Metadata written at the start of a Broadcast Wave file works with most systems, but some expect metadata to come at the end of the file. For these systems, choose the option to write metadata At the End of the File.

UMID Specifies Unique Material Identifier Data (UMID) according to the SPMTE 330M Standard. This information is read-only.

Using Shift-JIS Enables display and editing of Japanese characters in the existing text fields, provided that you already have a Japanese font installed on your computer. If a file contains Shift-JIS metadata that is not displayed in the Broadcast Wave tab, the metadata will be preserved.



To insert a Broadcast Wave file in Multitrack View, see “Spot insert a Broadcast Wave file into a session” on page 47.

Cart options

Use the Cart tab of the File Info 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 File Info 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.

Add or edit XMP metadata

The XMP File Info command lets you add Extensible Metadata Platform (XMP) information to audio files that are in WAV, mp3, or CEL format. Applications that support XMP—including all Adobe applications—can read, edit, and share this XMP metadata across databases, file formats, and platforms.

XMP metadata gives you many more information fields than File Info metadata.

- 1 In Edit View, choose File > XMP Info.
- 2 Click items the list on the left side of the dialog box to access different sets of properties.
- 3 Set properties as desired, and click OK.

Note: In Adobe Bridge, you can read XMP metadata embedded in a WAV or mp3 audio file either in the Metadata panel or by choosing File > File Info. (See “Importing with Adobe Bridge” on page 50.)

Chapter 15: Automating tasks

Favorites

Apply favorites to audio files

Favorites are effects, scripts, and even third-party tools that you can save for easy access. The Favorites panel lists some defaults to get you started, plus any additional favorites you create. The same items appear in the Favorites menu.

❖ In Edit View, double-click the favorite in the Favorites panel, or choose it from the Favorites menu.

See also

“About scripts” on page 261

“Applying effects in Edit View” on page 106

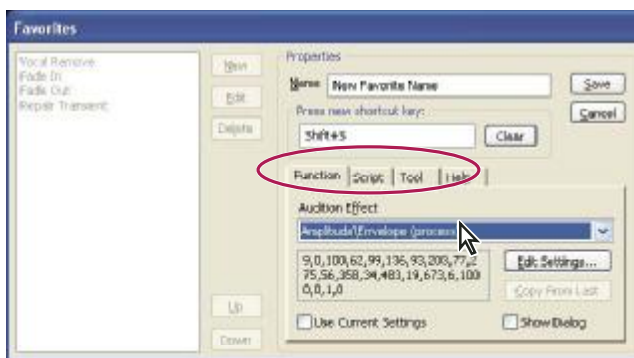
Create, edit, and organize favorites

1 In Edit View, choose Favorites > Edit Favorites.

2 Do any of the following:

- To create a new favorite, click New. If you want to base the new favorite on an existing favorite, select the existing favorite first.
- To edit a favorite, select the favorite, and click Edit.
- To delete a favorite, select the favorite, and click Delete.
- To change the order of favorites, select a favorite, and click Up or Down.

3 If you are creating or editing a favorite, set options on the right side of the dialog box. Click the Function, Script, or Tool tab to access specific favorite options.



Click tabs in the Favorites dialog box to access specific settings.

4 When you finish setting options, click Save, and then click Close.

Favorites options

You can set the following options when you create or edit a favorite:

Name Specifies the name of a favorite. To organize the Favorites menu, do one or more of the following:

- Create hierarchical menus by using a backslash (\). For example, type **My Delays\Hall Reverb** in the Name box to place the Hall Reverb favorite in the My Delays submenu.
- Add separator bars by typing a series of dashes (-----) into the Name 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 and typing three hyphens (such as **My Effects\---**). You can create only one separator bar on each submenu.

Note: If you name a new submenu, make sure not to use the name of any command, script, or tool that is listed under the Function, Script, or Tool tab.

Press New Shortcut Key Type here to assign a key or combination of keys as the keyboard shortcut for 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. Click the Clear button to delete an existing shortcut.

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.

Function 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 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, which 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 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, which lets you navigate to the desired tool.

Help Displays instructions for adding separators and submenus to the Favorites menu.

Normalizing groups of files

Normalize a group 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-stage batch process. If volume is raised during this 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.*

- 1 In Edit View or CD View, import the waveforms you want to normalize.
- 2 Choose Edit > Group Waveform Normalize.
- 3 Select the 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.
- 4 Click the Analyze Loudness tab, and then click Analyze 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 “Analyze Loudness statistics” on page 258.)
- 5 Click the Normalize tab, and specify how you want to normalize the waveforms. (See “Group Waveform Normalize options” on page 259.)
- 6 Click Run Normalize.

See also

“Normalize effect (Edit View only)” on page 120

Analyze Loudness statistics

In the Group Waveform Normalize dialog box, click the Analyze Loudness tab, and then click Analyze Statistical Information to display the following data:

Eq-Loud Is the average loudness value with an equal-loudness equalization curve that takes into account frequencies to which the human ear is most sensitive.


Loud Is the average loudness value without equal-loudness equalization.

Max Is the maximum RMS (Root-Mean-Square) amplitude present.

Avg Is the average RMS amplitude 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.*

See also

“Normalize a group of files” on page 258

Group Waveform Normalize options

In the Group Waveform Normalize dialog box, click the Normalize tab to display the following options:

Normalization Specifies whether to normalize to an average level or a specific level you enter in decibels.

Note: Unlike the Normalize effect, this Normalization option doesn't use percentages, 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 When you amplify audio, the audio samples may extend beyond the clipping point. If out-of-band peaks occur, you can choose to clip the waveform (causing distortion) or apply limiting (increasing perceived volume without distortion).

Options include the following:

- **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 option provides two additional options: Lookahead 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 below normal levels for a while.

Statistics RMS Width Specifies the length of audio segments used to calculate the Analyze Loudness statistics. If you change this setting, click Analyze Statistical Information to update the statistics.

See also

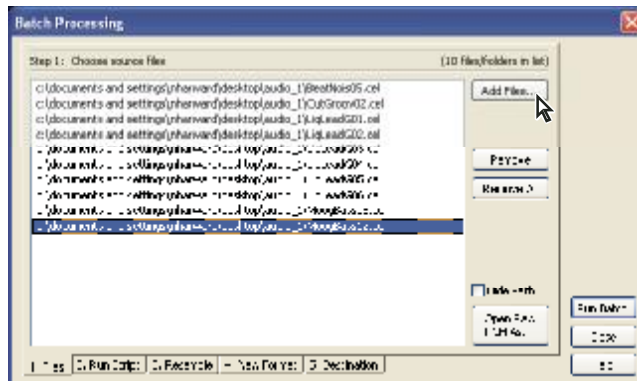
“Normalize a group of files” on page 258

Batch processing files

Batch processing overview

The Batch Processing dialog box in Adobe Audition enables you to process multiple files in three possible ways; you can run a script, resample, or save to a new file format.

Adobe Audition comes with several scripts that you can use to batch process files. You can also create your own scripts. To use a script for batch processing, you must record the script on an open waveform with no selection.



Adding files in the Batch Processing dialog box

See also

“About scripts” on page 261

Batch process files

- 1 In Edit View, choose File > Batch Processing.
- 2 Click the tabs at the bottom of the dialog box to set the batch processing options described below, and then click Run Batch.

Files Click Add Files, choose the source files you want to batch process, and click Add. (If the source files are Raw PCM files, click Open Raw PCM As to select the desired Sample Rate, Channels, Resolution, and other properties.)

Once you choose source files, you can use the Remove and Remove All buttons to delete files from the list. You can also select Hide Path to display the name of the file without its full path.

Run Script (Optional) Select Run A Script to run a script on the files. Click Browse to select a script collection (SCP) file, and click Open. Then choose a script from the Script menu.

If the script you want to run doesn't appear in the Script menu, it may be because you didn't record the script on an open waveform with no selection made. In this case, close the Batch Processing dialog box and rerecord the script.

Resample (Optional) 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.

For more information on options, see “Convert the sample rate of a file” on page 98.

New Format Choose an output format for the files.

Sample Format Types lists the sample properties of the waveforms that are to be converted. If more than one entry is listed, you may want to click Format Properties and set different properties for each. For instance, a 22 kHz mono waveform might need different encoding options than a 44 kHz stereo file.

Destination Select a destination folder and specify how to rename the files. (See “Batch file renaming options” on page 260.)

Batch file renaming options

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 file names 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 file name 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 file name or extension.

Here are some examples of how files can be renamed:

Original name	Output Filename Template name	Resulting file name
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

See also

“Batch process files” on page 260

Scripts

About scripts

In Edit View, 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 effect-specific settings, and then apply them at any time simply by running a script.

Adobe Audition lets you create three types of scripts:

- Scripts that start from scratch on a newly created waveform; their first command is File > New.
- Scripts that work on the currently open waveform, from the start-time indicator to the end of the file.
- Scripts that work on a selected range in an open waveform.


A set of scripts can be grouped together in a *script collection*. For example, a script collection called “ambience” might contain scripts for adding echo, reverb, and delay, and one called “batch utilities” might contain scripts for batch processing.

See also

“Batch processing overview” on page 259

Create a script

- 1 Set up Edit View for the script: If you’re creating a script that applies to open waveforms, open a file typical of the ones you’ll apply the script to. Then select a range if needed.
- 2 Choose File > Scripts. The Script Collections area displays the name of the currently opened script collection.
- 3 If the script collection you want isn’t open, 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 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.
- 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.

 *Don’t open or save a file as part of the actions for the script—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 finish recording the script, choose File > Scripts, and click Stop Current Script.

- 8 Type a description for the script in the text area of the dialog box.

The description appears when the script is selected. You can add or edit a description later by clicking Edit Script File.

- 9 Click Add to Collection. The script appears in the list at the left.

Run a script

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. See “Batch process files” on page 260.*

- 1 Set up Edit View 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 range. If you want to run a script that starts from scratch, close all open waveforms.
- 2 Choose File > Scripts. The Script Collections area displays the name of the currently opened script collection.
- 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, 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 script, performs all script operations relative to the original position of the start-time indicator, as opposed to the current position.

For example, if a script was recorded with the start-time indicator at 0:10:00, selecting this option applies the script at the current position, plus 10 seconds: If the current 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 start-time indicator position, record the script with the indicator at a 0:00:00 position, and select this option when you run it.

Script description area Shows any description entered for the script and indicates how the selected script works: Script Starts From Scratch creates a new file; Script Works On Current Wave works on an entire waveform; and Script Works On Highlighted Section works on a selected range of a waveform.

Edit a script

- 1 In Edit View, choose File > Scripts. The Script Collections area displays the name of the currently opened script collection.
- 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.

Chapter 16: Building audio CDs

Assembling CDs


About CD View

CD View provides an easy-to-use interface for assembling CD tracks, setting track properties, and burning CDs. You can assemble tracks for a CD all at once, or you can insert individual tracks as you finish editing audio. After you insert tracks, you can reorder or remove them.

When you assemble audio for a CD, you'll probably want to fine-tune individual tracks so they form a cohesive whole. This process—known as *mastering*—often involves cropping files, applying compression, and comparing tracks to ensure continuity of levels and EQ. For more information, see “Mastering effect” on page 161 and “Normalize a group of files” on page 258.



Compact discs typically hold 650 MB of data, which equals 74 minutes of stereo audio.



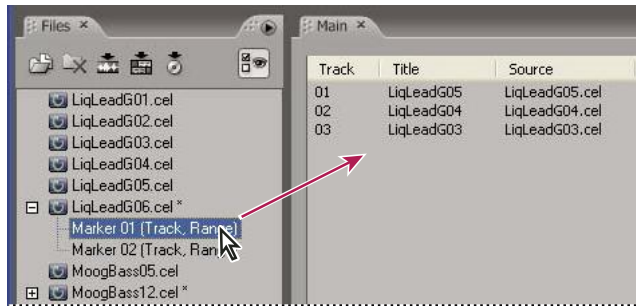
CD View

See also

“Basic components of Edit, Multitrack, and CD View” on page 19


Insert CD tracks

Adobe Audition provides a variety of ways to insert tracks into CD View. Keep in mind that you're not limited to inserting entire files—you can also insert audio ranges that are defined with CD track markers.



Dragging a marker to the track list

❖ Do any of the following:

- Select one or more files or track markers in the Files panel. Then drag the items into the track list in CD View or click the Insert Into CD List button  in the Files panel. If you insert a file that contains track markers, a dialog box appears asking if you want to insert the file as a single track or as multiple tracks based on the markers.

Note: If you insert a marker point rather than a range, you are prompted to select start and end points from a list of available markers. Adobe Audition then creates a new marker range and inserts it into the CD project.

- In CD View, choose Insert > Audio or Insert > Audio From Video. Select a file, and click Open.
- In CD View, choose Insert > File/Marker List. Select the file or track marker 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 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 Into CD List.
- In Multitrack View, open a session file, and choose File > Export > Audio Mix Down. Then select Insert Mixdown Into, choose CD View, and click Save. By default, any track marker ranges are automatically inserted as separate CD tracks. For more information, see “Export a session to an audio file” on page 242.



To divide one long audio file (such as a concert recording that includes several songs) into multiple CD tracks, insert the file into a session, and add track markers at the desired locations. Then export an audio mix down to CD View.

See also

“Files panel overview” on page 48

“About markers” on page 80

“Adjust, merge, convert, or delete markers” on page 82

Select CD tracks

❖ In CD View, do any of the following:

- To select a track, click the track in the track list.
- To select adjacent (contiguous) tracks, click the first track in the desired range, and then Shift-click the last.
- To select nonadjacent (noncontiguous) tracks, Ctrl-click them.
- To select all tracks, choose Edit > Select All Tracks.

Rearrange, reassign, or remove CD tracks

❖ In CD View, do any of the following:

- To rearrange tracks, drag the track you want to move up or down in the list. Or, select the track and click the Move Up or Move Down button.
- To remove tracks, select one or more tracks and click Remove.
- To remove all tracks, click Remove All.
- To close the source files when removing tracks, select one or more tracks and choose Edit > Destroy Selected Tracks (Remove And Close).
- To reassign a track to an available track marker or the entire file, right-click the track and choose Change Selection from the context menu. The submenu lists all available track marker ranges in the file.

Edit the source audio for a CD track

1 In CD View, select the track you want to edit.

2 Choose Edit > Edit Source Audio.

See also

“Editing audio files” on page 65

Set CD 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.

1 In CD View, double-click a track to open its track properties. Or, select the track and either click Track Properties or choose View > Track Properties.

2 Enter a track title and artist for the track.

Important: To add title and artist text to a CD, you must select Write CD-Text in the Write CD dialog box, and your CD burner must support this text. (See “Write a CD” on page 267.)

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 0-second pause to the beginning of every track except the first track, which must have a 2-second pause that cannot be changed.

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.

Set As Default Uses the current settings as the default track properties.

Same For All Tracks Applies the settings, with the exception of the ISRC code, to all tracks in the track list.

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 U.S.A.)
- Registrant code: 3-digit, alphanumeric, unique reference
- Year of reference: last 2 digits of the year (for example, 04 for 2004)
- Designation code: a 5-digit, unique number

Save and open CD lists

You can save tracks assembled in CD View as a *CD list*. A CD list contains references to all the tracks in your CD project and gives their order and track properties. Saving a CD list gives you the flexibility to continue work on a CD project at a later time. Adobe Audition saves CD lists in CDL file format.

❖ Do any of the following:

- To save a CD list, assemble your tracks in CD View and then choose File > Save CD List.
- To open a new, blank CD list, choose File > New CD List
- To open an existing CD list, choose File > Open CD List.

Writing CDs

Set CD device properties

Before you write a CD, you should verify that your CD burning device is set up correctly.

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.

- 1 In CD 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.

Write a CD

- 1 Insert a blank, writable CD into the CD burning device.
- 2 In CD View, click Write CD or choose File > Write CD.
- 3 Set the following options:

Device Specifies the device you want to use to write the CD. Click Device Properties to set device properties. (See “Set CD device properties” on page 267.)

Write Mode Specifies whether on not to test the CD for buffer underruns. Choose from the following options:

- **Write** Writes the CD without testing for buffer underruns.
- **Test** Tests if the CD can be written without the occurrence of buffer underruns. No audio is written to the CD.

- **Test and Write** Tests for buffer underruns and then proceeds with the actual write process if the test is successful.

Copies Specifies how many copies of the CD you want to burn. You can burn up to 99 copies of a CD. When burning multiple copies of a CD, Adobe Audition prompts you to insert a new CD for each copy.

Eject CD When Complete Ejects the CD tray upon completion of the write process.

Write CD-Text Specifies the Title, Artist name, and UPC/EAN code for the CD. The title and artist name you type here applies to the CD as a whole; it does not overwrite titles or artists for individual tracks. (See “Set CD track properties” on page 266.)

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.

- 4 Click Write CD. The Track and Disk bars show you the progress of the write process.


Chapter 17: Keyboard shortcuts

Finding and customizing shortcuts

About keyboard shortcuts and MIDI triggers

Adobe Audition provides a default set of keyboard shortcuts to help you speed up the editing process. In menus and tool tips, available keyboard shortcuts appear to the right of command and button names.

You can customize nearly all 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 called a *MIDI trigger*. (For example, you can assign the Play command in Adobe Audition to the C4 note on your MIDI keyboard.)

 *Adobe Audition also features a shortcut bar with groups of buttons you can show or hide to suit your current task. For more information, see “Display the shortcut bar” on page 27.*

Find shortcuts

❖ Do any of the following:

- For menu commands, look for shortcuts to the right of command names.
- For tools or buttons, look for shortcuts to the right of tool tips. (To display tool tips, hold the pointer over a tool or button.)
- For a complete list of all shortcuts, choose Edit > Keyboard Shortcuts And MIDI Triggers. For Category, select Show All.

Note: *Many shortcuts work only when a specific view is active. To show or hide shortcuts related to specific views, click the Multitrack, Edit View, and CD View buttons in the Keyboard Shortcuts And MIDI Triggers dialog box.*

Customize shortcuts

1 Choose Edit > Keyboard Shortcuts And MIDI Triggers.

2 Filter the list of commands as needed:

- Select (Show All) or a specific group from the Category menu.
- Click the Multitrack, Edit View, and CD View buttons to access commands for each application view.

3 In the Command Name column, select the command you want to customize.

4 Do any of the following:

- To assign a keyboard shortcut, click inside the Press New Shortcut Key box, and press the desired key combination.
- To assign a MIDI trigger, click inside the Press New MIDI Note box, and press the desired key on a MIDI keyboard, or adjust a MIDI controller (such as a foot pedal).
- To remove a keyboard shortcut or MIDI trigger, click the appropriate Clear button.

Note: *If you enter a key combination that’s already in use, Adobe Audition notifies you in the Conflicting Keys box. Click Clear, and enter a different shortcut before continuing.*

- 5 (Optional) To save a customized shortcut set, click Save As, type a unique name, and click Save.
- 6 When you finish making changes, click OK.

See also

“Enable MIDI-triggered shortcuts” on page 270

Restore the default shortcuts

- 1 Choose Edit > Keyboard Shortcuts And MIDI Triggers.
- 2 From the Set menu, choose Adobe Audition Default.

Enable MIDI-triggered shortcuts

- 1 Choose Edit > Keyboard Shortcuts And MIDI Triggers.
- 2 In the MIDI Trigger section, select the Channel and MIDI In port for the connected device. Then click OK.
- 3 In Edit View or Multitrack View, choose Options > MIDI Trigger Enable.

See also

“About keyboard shortcuts and MIDI triggers” on page 269

“Customize shortcuts” on page 269

Default keyboard shortcuts

Keys for opening views

The following keyboard shortcuts perform the same functions in Edit View, Multitrack View, and CD View. This partial list includes the shortcuts that many Adobe Audition experts find most useful.

Result	Shortcut
Open Edit View	8
Open Multitrack View	9
Open CD View	0
Toggle between Edit View and Multitrack View	F12

For a complete list of shortcuts, choose Edit > Keyboard Shortcuts And MIDI Triggers.

See also

“Customize shortcuts” on page 269

Keys for playing and zooming audio

The following keyboard shortcuts perform the same functions in either Edit View or Multitrack View. This partial list includes the shortcuts that many Adobe Audition experts find most useful.

Result	Shortcut
Start and stop playback	Spacebar
Select Scrub tool	A
Move start-time indicator to beginning of timeline	Home
Move start-time indicator to end of timeline	End
Zoom in horizontally	=
Zoom in vertically	Alt+=
Zoom out horizontally	-
Zoom out vertically	Alt+-

For a complete list of shortcuts, choose Edit > Keyboard Shortcuts And MIDI Triggers.

See also

“Customize shortcuts” on page 269

Keys for editing audio files

The following keyboard shortcuts apply only in Edit View. This partial list includes the shortcuts that many Adobe Audition experts find most useful.

Result	Shortcut
Open File Info dialog box	Ctrl+P
Repeat previous command (opening its dialog box but not clicking OK)	F2
Repeat previous command (opening its dialog box and clicking OK)	F3
Open Convert Sample Type dialog box	F11
Capture a noise reduction profile for the Noise Reduction effect	Alt+N
Play preroll and postroll (skip selection)	Shift+R
Activate left channel of a stereo file for editing	Up Arrow
Activate right channel of a stereo file for editing	Down Arrow

For a complete list of shortcuts, choose Edit > Keyboard Shortcuts And MIDI Triggers.

See also

“Customize shortcuts” on page 269

Keys for mixing multitrack sessions

The following keyboard shortcuts apply only in Multitrack View. This partial list includes the shortcuts that many Adobe Audition experts find most useful.

Result	Shortcut
Select the same input or output for all audio tracks	Ctrl+Shift-select
Activate or deactivate Mute, Solo, or Record in all tracks	Ctrl+Shift-click
Adjust knobs in 10-unit increments	Shift-drag
Adjust knobs in fractional increments of one-tenth	Ctrl-drag
Select Hybrid tool	R
Select Time Selection tool	S
Select Move/Copy Clip tool	V
Nudge selected clip to the left	Alt+Left Arrow
Nudge selected clip to the right	Alt+Right Arrow
Select next clip in active track	Ctrl+Down Arrow
Select previous clip in active track	Ctrl+Up Arrow
Move automation points by the same amount	Alt-drag
Move automation points by the same percentage	Ctrl-drag
Export an audio mix down	Ctrl+Shift+Alt+M

For a complete list of shortcuts, choose Edit > Keyboard Shortcuts And MIDI Triggers.

See also

“Customize shortcuts” on page 269

Chapter 18: Digital audio glossary

The glossary is your guide to unfamiliar terms in common audio workflows and multiple Adobe Audition features. If you don't find a term here, search for it in Help to find a feature-specific definition.

Common audio terms

A

ADAT A digital 8-track tape deck manufactured by Alesis Corporation that is very popular in recording studios.

ADC (analog-to-digital converter) The hardware that converts an analog audio or video signal into a digital signal that you can process with a computer.

aliasing Noise that occurs when a high frequency sound exceeds the Nyquist Frequency for a given sample rate. (See “Nyquist frequency” on page 277.) Most analog-to-digital converters prevent aliasing by filtering out sounds above the Nyquist Frequency.

amplitude Amplitude represents the volume 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 274.)

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.

ASIO (Audio Stream In/Out) A standard for low-latency drivers, created by Steinberg Media Technologies.

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.

automation The process of recording volume, pan, and effects changes during a mix, and perfectly reproducing those changes every time a mix plays. (See “Techniques for automating mixes” on page 203.)

B

band pass filter A filter that allows some audio frequencies to pass through unchanged.

beats per minute (bpm) Musical tempo, which is defined by the number of beats that occur every 60 seconds.

bit depth (or bit resolution) The number of bits used to represent audio amplitude. For example, 8-bit resolution provides 256 possible amplitude levels and a 48 dB dynamic range; 16-bit resolution provides 65,536 levels and a 96 dB range. Adobe Audition supports up to 32-bit resolution with 4,294,967,296 possible levels. For the best audio quality, remain at 32-bit resolution while transforming audio in Adobe Audition, and then convert to a lower bit depth 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.

C

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 depth (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 for the data compression schemes used by the ACM, AVI, MPEG, and QuickTime formats and the analog-to-digital converters on some sound cards. (Note that codecs only compress file size; to compress audio amplitude, apply a compressor effect.)

compressor An effect that reduces dynamic range by lowering amplitude when an audio signal rises above a specified threshold. For example, a compressor can compensate for variations in level caused by a vocalist who occasionally moves away from a microphone. Or, during mastering, a compressor can produce consistent levels for full program material, providing a solid, professional sound for web, video, and radio presentations. (See "Mastering effect" on page 161.) Adobe Audition provides two compressor effects: Dynamics Processing and Multiband Compressor.

crossfade A fade from one audio clip or 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.

D

DAC (digital-to-analog converter) The hardware that converts 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*.)

DAW (digital audio workstation) A computer system used to edit, process, or mix audio.

dBFS Decibels below full scale in digital audio. The maximum possible amplitude is 0 dBFS; all amplitudes below that are expressed as negative numbers. A given dBFS value does not directly correspond to the original sound pressure level measured in acoustic dB.

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.

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 audio file. In Adobe Audition, Edit View is a destructive editing environment; however, edits do not permanently change audio until you save a file.

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.

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 Describes an audio signal without any signal processing such as reverb; the opposite of *wet*.

DSP (digital signal processing) The process of transforming a digital audio signal by using complex algorithms. Examples include filtering with equalizers, and effects processing with reverbs and delays.

DVD A storage medium similar to a CD, but with much higher bandwidth and storage capabilities. Audio in DVD movies generally uses a 96-kHz sample rate and a 24 bit depth.

dynamic range Audio amplitude range, from quietest troughs to loudest peaks. See also “bit depth (or bit resolution)” on page 273 and “decibel (dB)” on page 274.

E

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

FFT (Fast Fourier Transform) 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 sine and cosine functions, allowing frequency and amplitude to be quickly analyzed. Higher FFT sizes create more precise results but take longer to process.

frequency Describes the rate at which a sound wave vibrates, measured in cycles per second, or hertz (Hz). A cycle consists of a single, repeated sequence of pressure changes, from zero pressure, to high pressure, to low pressure, and back to zero. A sound wave's frequency determines its pitch: high frequency equals high pitch, and low frequency equals low pitch. (See also “Waveform measurements” on page 8.)

H

hertz (Hz) Cycles per second. A unit of measurement that describes the frequency of a sound. (See “frequency” on page 275.)

I

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

latency Measures the delay between user input and sound output from a computer. If latency is high, it produces an audible echo during recording, disrupting timing for musicians. To reduce latency, use sound cards with ASIO drivers.

limiter A signal processor that prevents audio from clipping. If the input signal exceeds the specified threshold level, 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. (See “About loops” on page 221.)

M

mastering The process of finalizing audio for a specific medium, such as the web or an 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. (See “Mastering effect” on page 161.)

MIDI Musical Instrument Digital Interface, a method of communicating performance instructions 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 between 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.

mix (or mix down) To combine multiple audio sources or tracks and output them together. Though mixes are typically output to a stereo pair of channels, they can be directed to any number of channels (for example, one channel for mono, or six channels for surround-sound).

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. (See also “dithering” on page 275.)

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 so it nearly reaches 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 A frequency equal to half the current sample rate, which determines the highest reproducible audio frequency for that rate. For example, audio CDs use a sample rate of 44,100 Hz because the resulting Nyquist frequency is 22,050 Hz—just above the limit of human hearing, 20,000 Hz. For the best audio quality, record and edit at higher sample rates and then convert down if needed.

O

offline processing Intensive effects processing that requires dedicated computer power, briefly preventing you from editing audio. (Compare with “real time” on page 277.)

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.

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. (See also “Waveform measurements” on page 8.)

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.

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.

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 later 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 depth. This noise is highest at low bit depths, 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 immediately respond to user input. Note, however, that system speed ultimately determines processing time. Adobe Audition provides real-time mixing and effects in Multitrack View, and real-time previews in Edit View. (Compare with “offline processing” on page 277.)

resample To convert a sound file to a different sample rate and bit depth.

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 brain can't separate them. Adobe Audition offers three reverb effects: Studio Reverb, Reverb, and Full Reverb.

ReWire A technology for synchronizing audio applications created by Propellerhead Software.

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 277.)

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 piece of software or hardware 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 276.)

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.

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.

spectral editing Editing audio using a display that visually represents audio frequencies. In spectral displays, bass frequencies appear at the bottom, treble at the top. Volume is represented by color intensity. In Adobe Audition, you perform spectral editing to remove noise or process specific frequency ranges.

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 beats per minute. (See "beats per minute (bpm)" on page 273.)

timecode An audio or digital signal that synchronizes time between multiple devices. The most common forms are SMPTE and MIDI timecode.

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.)

V

VST Virtual Studio Technology, a plug-in format compatible with a wide variety of audio software. VST plug-ins provide audio effects such as compression and reverb; VSTi plug-ins provide virtual instruments such as samplers and synthesizers.

W

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.)

wet Describes 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.

Index

Numerics

- 5.1 surround sound, about 232
- 64-bit doubles (RAW) format 244
- 8-bit signed format 244

A

- ACM Waveform format 245
- activation of software 1
- Adaptive Noise Reduction effect 141
- ADC, defined 273
- Add Tracks command 182
- Add/Remove VST Directory command 110
- Adjust Boundaries command 195
- Adjust Sample Rate command 98
- Adobe After Effects
 - editing audio clips 227
 - remixing soundtracks 227
- Adobe Audition
 - and analog-to-digital conversion 12
 - Loopology DVD 2
 - new features in version 3.0 5
 - preferences 33
 - workflow overview 13
- Adobe Bridge
 - about 50
 - XMP metadata 255
- Adobe Design Center 3
- Adobe Help 2
- Adobe Premiere Pro, remixing soundtracks 227
- ADPCM compression 246, 247
- Advanced Session Properties dialog box 210
- After Effects. *See* Adobe After Effects
- AIFF format, about 245
- A-Law Wave format 244
- aliasing, defined 273
- Always Input command 57
- Amiga IFF-8SVX format 245
- Amplify effect 112
- Amplify/Fade effect 112

- amplitude
 - defined 273
 - visually adjusting 80
- Amplitude Statistics dialog box 96, 97
- analog audio, fundamentals of 10
- Analog Delay effect 122
- analog recording, defined 273
- analyzing
 - amplitude statistics 96
 - frequency 93, 94
 - loudness 258
 - phase 91
- Apple AIFF format 245
- ASCII text data 245
- ASIO, defined 273
- ASPI/SPTI options 35, 52
- attack, defined 273
- attenuate, defined 273
- AU format, about 249
- audio
 - See also* playing audio, recording audio
 - analyzing 96
 - displaying 65, 71
 - fundamentals of 8
 - inverting 87
 - mastering 161
 - muting 86
 - restoring 141
 - shrinking or stretching 172, 200
- Audio Clip Properties window 202
- audio files
 - See also specific format names*
 - adding file information 253, 255
 - appending to 44
 - closing 49, 244
 - creating 43
 - editing permanently 20
 - exporting mix downs to 242
 - formats, selecting 244
 - inserting in sessions 47
 - saving marked ranges 83
 - saving to one format 240
 - saving to various formats 240
 - size of 11, 28

- Audio Mix Down command 242
- Audition Loop format 246
- Audition. *See* Adobe Audition
- Auto Heal 74
- Auto Play
 - in Files panel 50
 - in Open dialog box 44
 - preferences 33
- Auto-Mark 84
- Automatic Click Remover effect 142
- Automatic Phase Correction effect 142
- automating tasks
 - batch processing files 259
 - normalizing groups of files 258
 - using favorites 256
- automation
 - about 203
 - clip envelopes, using 203
 - defined 273
 - editing 207
 - lanes 204
 - recording 205
 - track automation modes 205
 - track envelopes, using 204
- AVI format
 - exporting mix downs to 243
 - uncompressed 252
- azimuth misalignment, correcting 142

B

- band pass filter, defined 273
- Bars And Beats time format 55, 221, 225
- bass management circuit 236
- batch processing files 259, 260
- beats per minute (bpm)
 - and tempo 225
 - defined 273
- beats, finding 75
- Binaural Auto-Panner effect 162

- bit depth
 - automatically converting files to 32-bit 37
 - changing 99
 - defined 273
 - fundamentals of 11
 - bit resolution, defined 273
 - bitmap (BMP) images
 - importing as audio 44
 - saving audio as 250
 - blank files, creating 43
 - Bounce To New Track command 201
 - brightness of interface 29
 - Bring To Front command 195
 - Broadcast Wave format
 - options 254
 - preserving metadata 240
 - spot-inserting files 47
 - brown noise
 - about 89
 - defined 273
 - burning CDs. *See* CDs
 - buses
 - about 186
 - adding to sessions 182
 - converting to clips 201
 - defined 274
 - BWF format 47, 240, 254
- C**
- Cart Chunk information 255
 - CBR encoding 238, 248, 250
 - CCITT formats 244, 245, 249
 - CD View
 - default workspace 29
 - switching to 20, 270
 - using 264
 - CDs
 - adding text 266, 268
 - buffer and write speed for burners 267
 - burning, using Adobe Audition 267
 - burning, using another program 240
 - copy protection 266
 - editing source audio 266
 - extracting tracks from 50
 - inserting tracks 264
 - normalizing volume 258
 - pre-emphasis 266
 - recording from 53
 - removing tracks 266
 - saving CD lists 267
 - selecting tracks 265
 - setting pauses between tracks 266
 - track properties 266
 - CEL format 246
 - Center Channel Extractor effect 163
 - center lines, showing or hiding 36
 - Channel Mixer effect 164
 - channels
 - converting between mono and stereo 98
 - editing individually 74
 - selecting data in left or right 74
 - Check For Hidden Clips command 195
 - Chorus effect 137
 - click track, defined 274
 - Click/Pop Eliminator effect 143
 - clip indicators 62, 63
 - Clip Restoration effect 147
 - clipboard
 - choosing 76
 - clearing on exit 35
 - resampling data 38
 - clipping
 - amplitude and 62
 - defined 274
 - showing or hiding clipping boundaries 37
 - clips
 - about 191
 - aligning 193
 - combining 201
 - copying 193
 - editing 195
 - envelopes for 203, 210
 - fading and crossfading 197
 - grouping 192
 - hidden 195
 - inserting in sessions 47
 - locking in time 202
 - moving 192
 - muting 202
 - naming 202
 - properties, changing 202
 - protecting from recording 202
 - recording 57
 - removing 195
 - repeating 194
 - selecting 192
 - snapping 193
 - splitting and rejoining 197
 - time stretching 200
 - video 228
 - closing files 49, 244
 - CODEC, defined 274
 - color preferences 29, 35
 - command line, playing audio from 62
 - compressor, defined 274
 - Constant Bit Rate. *See* CBR encoding
 - converting
 - audio file formats 240
 - bit depth 99
 - sample rates 98, 101, 242
 - sample type when opening 46
 - stereo and mono waveforms 98
 - Convolution effect 158
 - Convolution Reverb effect 153
 - copy protection 266
 - Copy Reference Here command 193
 - Copy To New command 224
 - Copy Unique Here command 193
 - Creative Sound Blaster format 246
 - crossfading
 - clips 197, 199
 - defined 274
 - crossover frequency, for surround subwoofer 232
 - crosstalk, defined 274
 - cues. *See* markers
 - current-time indicator. *See* start-time indicator
 - custom time format 55
 - customizing
 - keyboard shortcuts and MIDI triggers 269
 - spectral displays 69
 - timecode display 34
- D**
- DAC, defined 274
 - DAT, defined 274
 - Data preferences 37
 - DAW, defined 274
 - dB, defined 274
 - dBFS, defined 274

- DC offset
 - adjusting meters for 63
 - correcting with Amplify/Fade effect 113
 - correcting with Normalize effect 121
 - defined 274
 - viewing 97
 - decibel, defined 274
 - default settings
 - See also* preferences
 - default multitrack session 47
 - keyboard shortcut set 269
 - workspaces 29
 - Delay effect 123
 - delay, compared to echo 122
 - deleting
 - audio 78
 - silence 86
 - Design Center 3
 - Destroy Clips command 195
 - Destroy Selected Tracks (Remove And Close) command 266
 - destructive editing
 - defined 275
 - Edit View and 20
 - devices. *See* input devices, output devices, MIDI sequencing
 - Dialogic ADPCM format 246
 - DiamondWare Digitized format 247
 - digital audio
 - See also* audio
 - extraction 50
 - fundamentals of 10
 - Digital Theater Systems (DTS) 232
 - direct-to-file recording 57
 - DirectX plug-ins 110
 - disk space
 - freeing up 41
 - monitoring 28
 - Display preferences 36
 - Display Time Format command 55
 - Distortion effect 159
 - dithering
 - defined 275
 - enabling for WAV files 251
 - preferences 37
 - when lowering bit depth 38, 99
 - docking panels 23
 - Dolby Digital 232
 - Doppler Shifter effect 164
 - downloads
 - updates, plug-ins, and tryouts 4
 - downsampling quality level 38
 - dry, defined 275
 - DSP, defined 275
 - DTMF signals, generating 88
 - DTS 232
 - dubbing 57
 - Duplicate (clip) command 194
 - Duplicate Selected Track command 184
 - DV AVI format 252
 - DVD, defined 275
 - DVI/IMA ADPCM format 247
 - DWD format 247
 - Dynamic Delay effect 123
 - Dynamic EQ effect 128
 - Dynamic Peaks option 63
 - dynamic range
 - defined 275
 - in level meters 63
 - Dynamics Processing effect 113
- E**
- Echo Chamber effect 125
 - Echo effect 124
 - echo, compared to delay 122
 - Edit Favorites command 256
 - Edit In Adobe Audition command 227
 - Edit Original data, embedding in files 227
 - Edit Tempo command 225
 - Edit View
 - about 20
 - customizing spectral displays 69
 - default workspace 29
 - display modes 65
 - selecting time ranges 72
 - Spectral Frequency Display 66
 - Spectral Pan Display 66
 - Spectral Phase Display 67
 - switching to 20, 270
 - Waveform Display 65
 - workflow overview 13
 - editing audio
 - See also specific editing tasks*
 - destructive and nondestructive 20
 - in Edit View 65
 - in Multitrack View 195
 - in specific stereo channels 74
 - undoing and redoing 101
 - effects
 - applying in Edit View 106, 107
 - applying in Multitrack View 107
 - automating 204, 206
 - cancelling 38
 - freezing 109
 - graph controls 104
 - grouping in Effects panel 106
 - plug-ins 110
 - pre- and post-fader 109
 - presets 104
 - saving as favorites 256
 - Effects Paintbrush tool 73
 - Effects panel 106
 - Effects Rack 102, 107
 - empty audio clips, inserting 194
 - Enable DirectX Effects command 110
 - Enable Preroll And Postroll Preview command 107
 - encoding, lossless 250
 - Envelope effect 116
 - Envelope Follower effect 174
 - envelopes, mix automation
 - about 203
 - editing 207
 - for clips 203, 210
 - for tracks 204
 - equalization (EQ), defined 275
 - expander, defined 275
 - exporting
 - mix downs to audio 242
 - mix downs to video 243
 - Extensible Metadata platform (XMP) 253
 - External Controllers preferences 40
 - External monitoring 58
 - Extract Audio From CD command 51
 - extras 2
- F**
- fading
 - audio 79, 112
 - clips 197
 - fast-forwarding 55

- favorites
 - applying 256
 - creating, editing, and organizing 256
 - setting options 257
- FFT Filter effect 129
- FFT, defined 275
- file information
 - for loops 225
 - standard metadata 253
 - XMP 255
- files
 - See also* audio files, *specific format names*
 - assigning a bitmap image to 254
 - backwards compatibility 37
 - batch processing 259
 - opening in other applications 42
 - properties, adding 253, 255
 - saving audio to 240
 - size and contents of 11
- Files panel
 - about 48
 - display options 49
 - previewing audio in 49
 - using 49
- Find Beats And Mark command 224
- Find Beats commands 75
- fixing audio 141
- Flanger effect 139
- floating windows 23
- Flush Virtual File command 42
- flushing, forcing 35
- font installation 1
- frames, snapping to 76
- frequencies
 - analyzing 93, 95
 - defined 275
 - selecting 73
 - viewing 65
- Frequency Analysis panel 93, 94, 95
- Frequency Band Splitter effect 176
- Frequency Space Editing workspace 29
- Full command 195
- Full Reverb effect 153
- G**
- gain, for tracks 186
- General preferences 33
- graph controls 104
- Graphic Equalizer effect 131
- Graphic Panner effect 165
- grid lines, showing or hiding 36
- Group Waveform Normalize command 258
- Guitar Suite effect 160
- H**
- Hard Limiting effect
 - process version 118
 - VST version 117
- Help
 - about 2
- hertz (Hz), defined 275
- hidden clips, revealing 195
- Hiss Reduction effect 147
- horizontal scroll bar 21
- Hybrid tool 27, 181, 191
- I**
- IFF format, about 245
- images
 - importing as audio 44
 - saving audio as 250
- IMP files 158
- importing
 - audio files 47
 - MIDI files 216
 - video files 228
- impulse files
 - for Convolution effect 158
 - for Convolution Reverb effect 153
 - for Full Reverb effect 153
- input devices
 - about 31
 - adjusting levels 63, 64
 - choosing 31
- inputs, assigning to tracks 186
- Insert In CD List command 265
- Insert/Delete Time command 181
- inserting in sessions
 - audio files 47
 - empty audio clips 194
 - files from Files panel 49
 - MIDI files 214
 - video files 228
- interpolate, defined 275
- inverting audio 87
- ISRC number 266
- J**
- Jitter command 219
- K**
- karaoke mixes, creating 163
- key
 - setting for loops 222, 225
 - setting for sessions 223
- keyboard shortcuts
 - about 269
 - customizing 269
 - default 269, 270
 - for editing audio files 271
 - for mixing multitrack sessions 271
 - for playing and zooming audio 270
 - for switching views 270
- L**
- Lasso Selection tool 73
- latency, defined 276
- left channel, editing 74
- Levels panel 62, 63
- levels, adjusting 64
- LFE 232
- limiter, defined 276
- LiveDocs 2
- Loop Info tab 225
- Loop Properties command 222
- looping
 - during playback 60
 - while recording 58
- Loopology DVD 2
- loops
 - about 221
 - Adobe Audition Loopology DVD 221
 - creating 71, 224
 - defined 276
 - extending or shortening 222
 - key of 210, 223, 225
 - previewing 49
 - properties, overriding 222
 - properties, setting 225
 - synchronizing 221
 - tempo of 223, 225
- lossless encoding 250

M

Mackie Control 40

magnifying. *See* zooming

Main panel
 in Edit View 65
 in Multitrack View 179
 navigating in 20

Manage Temporary Folder Reserve Space command 41

markers
 about 80
 adding 82
 adding silence between 83
 automatically setting 84
 changing durations of 82
 deleting 82
 describing 82
 inserting in sessions 47
 merging and converting 82
 renaming 82
 saving 240
 selecting 82
 showing or hiding lines for 36
 snapping to 76
 types of 81

Markers panel 81

Marquee Selection tool 73

Master track
 about 186
 converting to a clip 201

mastering
 CDs 264
 defined 276
 with Mastering effect 161

Mastering And Analysis workspace 29

Mastering Rack 102, 106

Maximum Session (Dual Monitor) workspace 29

Maximum Waveform Editing (Dual Monitor) workspace 29

Merge/Rejoin Split command 197

merging markers 82

metadata
 adding to audio files 253
 XMP 255

meters
 about 62
 customizing 63
 monitoring levels with 63

metronome
 enabling 223
 setting up 211

Microsoft ACM format 245

Microsoft ADPCM format 247

Microsoft DirectX 237

MIDI
 defined 276
 file size 11
 using with Adobe Audition 32

MIDI Panic Button command 217

MIDI sequencing
See also VST instruments
 audio clips, converting to 201
 controller data, editing 219
 exporting and importing files 216
 fundamentals 213
 grid options 217
 input devices 215
 locking tracks 220
 note data, editing 218
 output devices 215, 216
 pitch, transposing 219
 presets 215
 recording 215
 selecting data 218
 snapping to musical scales 218
 tempo and time signature, setting 223
 tracks, inserting and adding 214
 velocity data, editing 219
 volume and pan settings 214
 zooming and scrolling tracks 217

MIDI Timecode (MTC), defined 276

MIDI triggers
 about 269
 customizing 269
 default 269, 270
 defined 276
 enabling 270

millisecond (ms), defined 276

miniplug, defined 276

mix downs, exporting
 to audio 201, 242
 to CD 265
 to video 243

Mix Paste command 77, 99

Mixer panel 179

mixing
 about 179
 automating 203
 default bit-resolution 39
 defined 276

m-Law Wave format 244

Monitor Input Level command 63

monitoring inputs
 with effects 57
 without effects 58

monitoring levels 63

mono waveforms, converting to stereo 98

mono, defined 276

mouse wheel, zooming with 34, 217

MOV format 252

Move/Copy Clip tool 191

mp3 format
 ID3 Tag 253
 options 247
 XMP metadata 255

MPEG format, importing audio from 43, 47

Multiband Compressor effect 118

multichannel WMA 238

multiple monitors 26

Multitap Delay effect 126

Multitrack preferences 38

Multitrack View
 about 20
 default workspace 29
 recording in 57
 switching to 20, 270
 work area of 179
 workflow overview 15

muting
 audio files 86
 clips 202
 tracks 184

N

naming clips 202

new features in Adobe Audition 3.0 5

New Session command 46

NeXT/Sun format 249

noise
 gate, defined 276
 generating 89

- shaping, defined 276
 - techniques for removing 141
 - Noise Reduction effect 149
 - nondestructive editing
 - defined 276
 - Multitrack View and 20
 - Normalize effect 120
 - normalizing
 - defined 277
 - files for CDs 264
 - groups of files 258, 259
 - Notch filter effect 132
 - notes for sessions, adding 212
 - Nyquist Frequency, defined 277
- O**
- offline editing, defined 277
 - OGG Vorbis format 249
 - on-clip fade and gain controls 78, 197
 - online resources 2
 - Open Append command 44
 - Open As command 46
 - Open Session command 46
 - opening
 - See also* inserting in sessions
 - CD lists 267
 - files in Edit View 43
 - files in other applications 42
 - order, defined 277
 - out-of-band peaks 259
 - output devices
 - adjusting playback levels 64
 - choosing 31
 - for burning CDs 267
 - outputs, assigning to tracks 186
 - overdrive, creating with Distortion effect 159
 - overdubbing 57
- P**
- Pan/Expand effect 166
 - panels
 - See also individual panel names*
 - docking and grouping 23
 - opening and closing 26
 - resizing 25
 - panning
 - preferences for individual sessions 210
 - preferences for Multitrack View 38
 - tracks in sessions 189, 214
 - Parametric Equalizer effect 132
 - pasting
 - audio in Edit View 77
 - clips in Multitrack View 193
 - resampling data when 38
 - pauses, adding between CD tracks 266
 - pausing playback or recording 55
 - PCM Raw Data format 251
 - PCM, defined 277
 - peak (PK) files, setting options for 37
 - peak indicators 62
 - phase
 - analyzing 91, 95
 - defined 277
 - illustrated 9
 - inverting for tracks 186
 - Phase Analysis panel, customizing display 92
 - pink noise
 - defined 277
 - generating 89
 - Pitch Bender effect 169
 - Pitch Correction effect 169
 - Pitch Shifter effect 171
 - Play Hidden Clips command 195
 - play lists 85
 - Play Looped button 60, 224
 - playback cursor 54
 - playback devices. *See* output devices
 - playing audio
 - automatically 44, 50
 - choosing a device for 31
 - linearly with Transport panel 60
 - looped with Transport panel 60
 - start point, setting 54
 - using Files panel 49
 - using Scrub tool 61
 - using Windows Run command 62
 - plug-ins
 - about 3
 - defined 277
 - effects 110
 - in Adobe Store 4
 - VST instruments 214
 - pre- and post-fader routing
 - effects 109
 - sends 188
 - pre-emphasis, for CDs 266
 - preferences
 - about 33
 - Colors 35
 - Data 37
 - Display 36
 - External Controllers 40
 - General 33
 - Multitrack 38
 - SMPTE/MTC 40
 - System 34
 - Premiere. *See* Adobe Premiere Pro
 - preroll and postroll
 - for effect previews 107
 - for selected ranges 61
 - options 61
 - presets
 - for effects 104
 - sample rate conversion 101
 - Sequencer 215
 - Spectral Controls 70
 - previewing
 - audio, in Files panel 49
 - effects, in Edit View 107
 - sample rate changes 98
 - process effects 104
 - punch-in
 - defined 277
 - recording 58
- Q**
- quality, for downsampling 38
 - quantization, defined 277
 - Quantize command 219
 - Quick Filter effect 134
 - QuickTime (MOV) format 252
- R**
- Radio Industry file data 253
 - Randomize Velocity command 219
 - ranges
 - saving to new file 240
 - selecting in Edit View 72
 - selecting in Multitrack View 181
 - RAW format 244, 251
 - RCA cable, defined 277

- read me file 1
 - real-time, defined 277
 - recording audio
 - See also* input devices
 - devices for, choosing 31
 - direct-to-file recording 57
 - in Edit View 56
 - in Multitrack View 57
 - punching-in 58
 - start point, setting 54
 - timed recording 56
 - redoing edits 101
 - referenced clips 193
 - registration of software 1
 - Remove All Tracks command 266
 - Remove Clips command 195
 - Remove Hidden Clips command 195
 - Remove Selected Tracks command 266
 - repeating commands 101
 - requirements, system 1
 - resample, defined 278
 - Rescale Volume Envelopes command 204
 - restoring audio 141
 - reverb
 - about 152
 - defined 278
 - effects 153, 156, 157
 - reversing audio 87
 - Revert To Saved command 101
 - rewinding 55
 - ReWire
 - about 208
 - converting to audio clips 201
 - defined 278
 - disabling 209
 - setting up 208
 - tracks, assigning devices to 209
 - right channel, editing 74
 - ripping, from CD 50
 - Ripple Delete command 195
 - RMS, defined 278
 - rulers
 - changing scale of 71
 - scrolling with 21
 - zooming with 21
- S**
- S/N ratio, defined 278
 - Safe During Write option 190, 206
 - SAM format 244
 - Sample Accurate Sync command 210
 - sample rate
 - conversion presets 101
 - converting for files 98
 - converting for sessions 242
 - defined 278
 - fundamentals of 10
 - previewing changes 98
 - sample type
 - converting when opening 46
 - displaying in status bar 28
 - sample, defined 278
 - sampler
 - defined 278
 - information, adding to files 253
 - SampleVision format 250
 - saving
 - audio between markers 83
 - audio files to one format 240
 - audio files to various formats 240
 - CD lists 267
 - copies of associated files 241
 - mix downs to audio 242
 - mix downs to video 243
 - sessions 241
 - Spectral Controls presets 70
 - Scientific Filters effect 135
 - scripts
 - about 261
 - collections 262
 - creating 262
 - editing 263
 - running on individual files 262
 - running on multiple files 259
 - saving as favorites 256
 - script collection (.scp) files 262
 - scrolling
 - preferences for Edit View 34
 - preferences for Multitrack View 39
 - using scroll bars or rulers 21
 - using Selection/View controls 22
 - Scrub tool 27, 61
 - Select All Tracks command 265
 - Select Entire Wave command 74
 - selecting audio
 - between beats 75
 - in Multitrack View 181
 - in one channel 74
 - using preroll and postroll 61
 - time ranges 72
 - tools for 27
 - at zero-crossing points 75
 - Selection/View Controls panel 22
 - sends
 - about 186
 - setting up 188
 - Sequencer panel
 - overview 213
 - See also* MIDI sequencing, VST instruments
 - sequencer, defined 278
 - Session Mixing workspace 29
 - Session Properties command 223
 - sessions
 - about 179
 - converting sample rates of 242
 - creating 46
 - default session 47
 - defined 278
 - deleting or inserting time in 181
 - inserting clips 47
 - linking to Adobe Premiere Pro or After Effects 227
 - notes about 212
 - opening 46
 - properties, advanced 210
 - recording in 57
 - saving 241
 - as SMPTE masters and slaves 209
 - Set Current Clipboard command 76
 - Settings command 33
 - shortcuts
 - See also* keyboard shortcuts
 - shortcut bar 27
 - showing and hiding
 - center lines 36
 - shortcut bar 27
 - status bar 27
 - toolbar (Tools panel) 27
 - shrinking audio 172, 200
 - silence
 - adding between markers 83
 - deleting 86

- generating 86
 - removing automatically 84
 - slip edit 196
 - Smart Input command 57
 - smoothing edit boundaries 37
 - SMP format 250
 - SMPTE
 - defined 278
 - masters and slaves 209
 - monitoring stability 28
 - SMPTE/MTC preferences 40
 - start time offset 209, 210
 - snapping
 - about 76
 - clips 193
 - loops 221
 - video frames 228
 - zero crossings 224
 - SND format 245, 249
 - software
 - activation 1
 - downloads 4
 - registration 1
 - software instruments. *See* VST instruments
 - solosing tracks 184
 - sound card, defined 278
 - sound wave, defined 278
 - sound, fundamentals of 8
 - source audio, editing in CD tracks 266
 - special effects
 - See also effects and specific effect names*
 - noise 89
 - tones 90
 - Spectral Controls panel
 - about 68
 - adjusting colors in 69
 - options in 70
 - saving presets 70
 - spectral displays
 - customizing 68, 69
 - selecting frequencies in 73
 - settings for 36
 - Spectral Frequency Display 66
 - Spectral Pan Display 66
 - Spectral Phase Display 67
 - switching 65
 - spectral editing, defined 278
 - splines for graphs 105
 - Split command 197
 - Spot Healing Brush 74
 - SPTI (SCSI Pass Through Interface) 35
 - SPTI options 52
 - start-time indicator 54
 - Static Peaks option 63
 - status bar 27
 - stereo
 - defined 278
 - imagery, changing 162
 - Stereo Expander effect 167
 - Stereo Field Rotate effect
 - process version 168
 - VST version 167
 - stereo waveforms
 - choosing a channel to edit 74
 - converting to mono 98
 - creating from mono 99
 - Stretch effect 172
 - stretching audio 172, 200
 - stripe, defined 278
 - Studio Reverb effect 157
 - subwoofer 234
 - Surround Encoder dialog box 232
 - surround mixes
 - about 232
 - adding tracks and buses 233
 - exporting 238
 - panning tracks 234, 235
 - previewing 236
 - sound cards for 237
 - Surround Panner 234
 - volume levels, adjusting 236
 - zooming waveforms for 236
 - SVX format 245
 - Sweeping Phaser effect 140
 - symmetric dithering 37
 - synchronizing
 - clips or cursor 54
 - loops 221
 - MIDI 32
 - ReWire 208
 - SMPTE 209
 - System preferences 34
 - system requirements 1
- ## T
- tempo
 - defined 278
 - editing 55
 - session tempo 211
 - temporary folders
 - managing size of 41
 - setting location 35
 - time display format 55
 - Time panel 54
 - Time Selection tool 181, 191
 - time stretching clips 200
 - timecode, defined 278
 - Timed Record mode 56
 - timeline, snapping to ruler 76
 - tones, generating 90
 - tool tips
 - for keyboard shortcuts 269
 - for shortcut bar 27
 - toolbar 27
 - See also* shortcut bar
 - tools
 - displaying 27
 - Effects Paintbrush 73
 - Hybrid 181, 191
 - Lasso Selection 73
 - Marquee Selection 73
 - Move/Copy Clip 191
 - Scrub 61
 - Spot Healing Brush 74
 - Time Selection 181, 191
 - Tools panel 27
 - Top/Tail view 71
 - touch tones, generating 88
 - track controls 179
 - Track EQ panel 190
 - tracks for CDs
 - editing source audio 266
 - inserting 264
 - removing 266
 - selecting 265
 - setting properties 266
 - tracks for sessions
 - adding 182
 - automating 204, 205
 - controls, showing and hiding 185
 - deleting 182
 - duplicating 184
 - equalizing 190

- gain for 186
 - identical settings, applying 184
 - inputs and outputs 186
 - inserting files into 49
 - moving 183
 - muting 184
 - naming 183
 - panning 189
 - phase, inverting 186
 - scrolling to 180
 - soling 184
 - summing to mono 189
 - types of 182
 - volume, setting 184
 - zooming 183
 - Transport panel
 - buttons in 55, 60
 - showing and hiding 53
 - transposing
 - MIDI notes 219
 - session key 223
 - triangles. *See* start-time indicator, markers
 - triggers. *See* MIDI triggers
 - Trim command 78, 195
 - tryouts 4
 - Tube-modeled Compressor effect 121
 - TXT format 245
- U**
- Uncompressed AVI format 252
 - undo
 - about 101
 - disabling 101
 - preferences for 34
 - unity gain, defined 279
 - UPC/EAN 268
 - updates 4
 - upsampling quality level 38
- V**
- Variable Bit Rate (VBR)
 - encoding 238, 248, 250
 - vertical ruler
 - changing scale of 71
 - zooming and scrolling 21
 - vertical scroll bar 21
 - vibrato, and Chorus effect 137
- video**
- clips 228
 - exporting mix downs to 230, 243
 - importing 228
 - previews, customizing 229
 - supported formats 252
 - thumbnail options 229
 - Video + Audio Session workspace 29
 - Video panel 229
 - views
 - See also* Edit View, Multitrack View, CD View
 - comparing 20
 - switching between 20
 - virtual instruments. *See* VST instruments ¶
 - VOC format 246
 - Vocoder effect 177
 - volume
 - making consistent 258
 - raising or lowering 80
 - VOX format 246
 - VST effects
 - defined 279
 - plug-ins, accessing 110
 - VST instruments
 - See also* MIDI sequencing
 - about 213
 - assigning and adjusting 214
 - installing and enabling 214
 - monitoring 214
- W**
- WAV format
 - ACM 245
 - DVI/IMA ADPCM 247
 - mu-Law 244
 - Windows PCM 251
 - XMP metadata 255
 - Wave Cache 34
 - Waveform Display
 - about 65
 - customizing 36
 - waveforms
 - defined 279
 - editing 65
 - fundamentals of 8
 - inverting 87
 - reversing 87
 - statistics 96
 - wet, defined 279
 - white noise
 - defined 279
 - generating 89
 - windows
 - floating 23
 - opening and closing 26
 - Windows Media Video (WMV) format 253
 - Windows PCM format 238, 251
 - Windows Recording Mixer command 64
 - WMA format, about 250
 - WMV format, about 253
 - work area
 - See also* workspaces
 - about 19
 - views 20
 - workflow
 - editing audio files 13
 - mixing multitrack sessions 15
 - overview 13
 - workspaces
 - about 23
 - changing colors 35
 - changing display settings 36
 - choosing 29
 - deleting 30
 - docking and grouping 23
 - drop zones 23
 - floating windows 23
 - multiple monitors 26
 - predefined 29
 - resetting 30
 - resizing panel groups 25
 - saving custom layouts 30
 - writing CDs 267
- X**
- XMP metadata 255
- Z**
- zero-crossing points
 - adjusting selections to 75
 - automatically adjusting edits to 39
 - defined 279
 - snapping to 76

zooming

- phase and frequency graphs 95
- using Selection/View controls 22
- using Zoom panel, scroll bars, or rulers 20