Using
ADOBE® SOUNDBOOTH® CS5

Contents

Chapter 10: Keyboard shortcuts

Chapter 11: Digital audio glossary

Chapter 1: What's new

Adobe® Soundbooth® CS5 includes more capabilities and content, giving you greater flexibility for audio editing and enhancement.

Improved launch experience Soundbooth automatically creates a three-track multitrack project, quickly getting you started. (See "Multitrack mixing and editing" on page 55.)

New sound effects and scores from Resource Central Resource Central includes 130 new scores and 6,000 new sound effects, bringing the total to 9,000.

Drag-and-drop from Resource Central Place downloaded sound effects and scores directly in the Editor panel. (See "Insert a score in a multitrack file" on page 62.)

Adjust track height Simply drag track dividers in multitrack documents. (See "Create, name, resize, or delete tracks" on page 56.)

Speech Search custom dictionaries Increase transcription accuracy by referencing text files. (See "Improve speech analysis with reference scripts" on page 23.)

Chapter 2: Digital audio fundamentals

An understanding of key audio concepts helps you get the most out of Adobe Soundbooth.

Understanding sound

Sound waves

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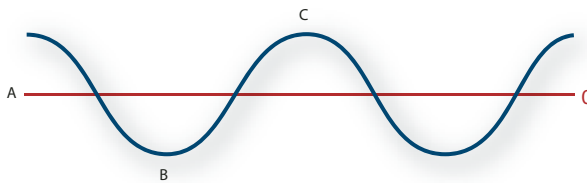
A sound wave represented as a visual waveform
A. Zero line B. Low-pressure area C. High-pressure area

Waveform measurements

Several measurements describe sound 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 90º at high pressure, 180º at the halfway point, 270º at low pressure, and 360º 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.
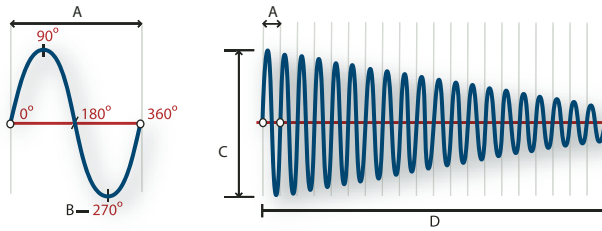
A single cycle at left; a complete, 20-Hz waveform at right
A. *Wavelength* **B.** *Degree of phase* **C.** *Amplitude* **D.** *One second*

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.
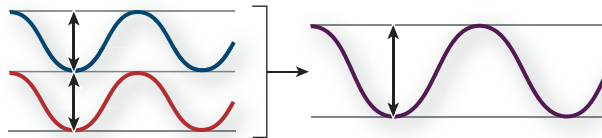
In-phase waves reinforce each other.

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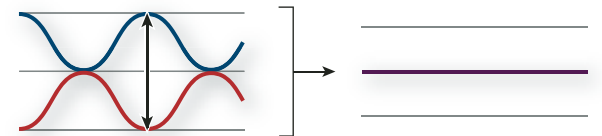
Out-of-phase waves cancel each other.

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.

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.
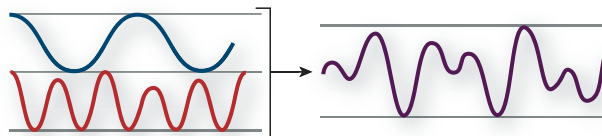
Two simple waves combine to create a complex wave.

Digitizing audio

Comparing analog and digital audio

In analog and digital audio, sound is transmitted and stored in very different ways.

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.

Understanding sample rate

Sample rate indicates the number of digital samples 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.
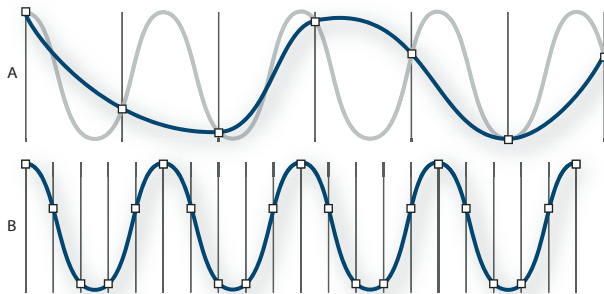
Two sample rates
*A. Low sample rate that distorts the original sound wave **B.** High sample rate that perfectly reproduces the original sound wave*

To reproduce a given frequency, the sample rate must be at least twice that frequency. (See "Nyquist frequency" on page 80.) 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 just beyond the limit of human hearing, 20,000 Hz.

The following table lists 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

Understanding bit depth

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. For the best audio quality, remain at 32-bit resolution while transforming audio in Soundbooth, and then convert to a lower bit depth for output.

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
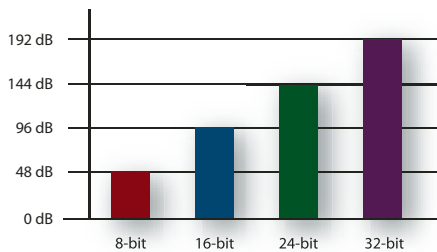
Higher bit depths provide greater dynamic range.

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

How Soundbooth digitizes audio

When you record audio in Soundbooth, 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. Soundbooth stores each sample in sequence until you stop recording.

When you play a file in Soundbooth, the process happens in reverse. Soundbooth 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 sum up, 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, Soundbooth can record, edit, and process it—the possibilities are limited only by your imagination.

Chapter 3: Workspace

Adobe Soundbooth provides a flexible workspace that you can quickly optimize for your working style.

Customizing the workspace

About workspaces

Adobe video and audio applications provide a consistent, customizable workspace. Although each application has its own set of panels (such as Project, Metadata, and Timeline), 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. As you rearrange panels, the other panels resize automatically to fit the window. You can create and save several custom workspaces for different tasks—for example, one for editing and one for previewing.

You can use floating windows to create a workspace more like workspaces in previous versions of Adobe applications, or to place panels on multiple monitors.
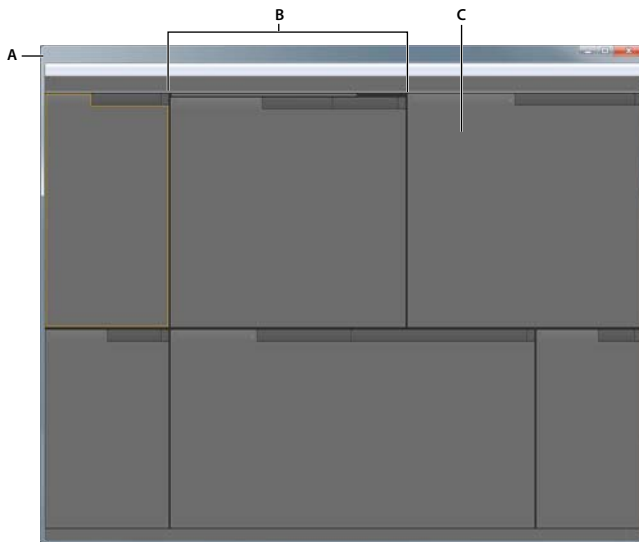
Example workspace
A. Application window B. Grouped panels C. Individual panel

Dock, group, or float panels

You can dock panels together, move them into or out of groups, and undock them so they float 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.
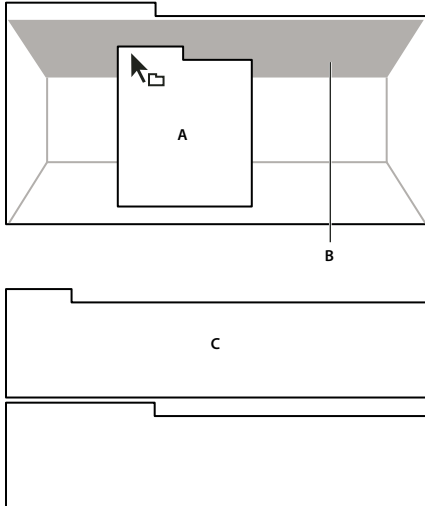
Dragging panel (A) onto docking zone (B) to dock it (C)

Grouping zones

Grouping zones exist in the middle of a panel or group, and along the tab area of panels. Dropping a panel on a grouping zone stacks it with other panels.
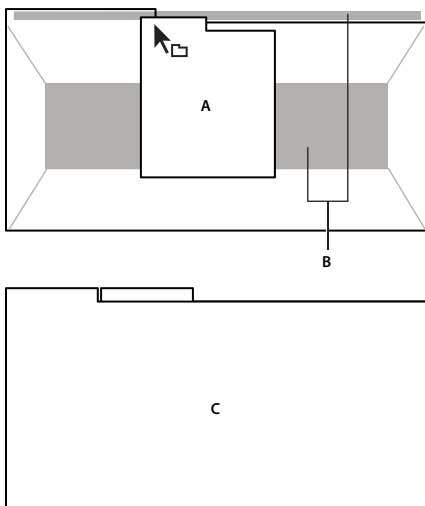
Dragging panel (A) onto grouping zone (B) to group it with existing panels (C)

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.
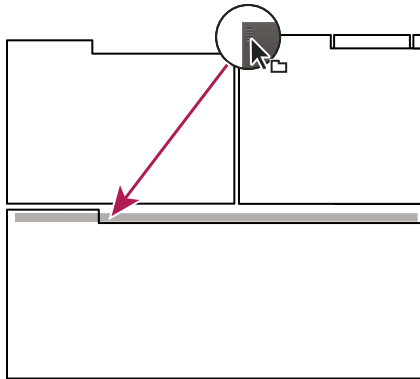
Drag panel gripper to move one panel

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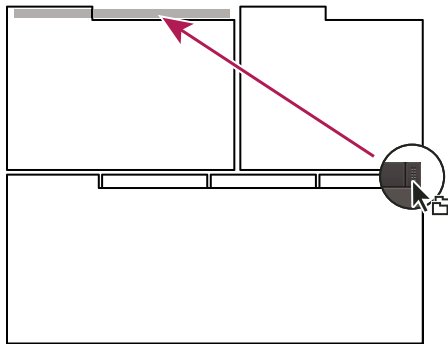
Drag group gripper to move entire group

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 and modify it similarly to the application window. You can use floating windows to use a secondary monitor, or to create workspaces like the workspaces 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 taskbar.)

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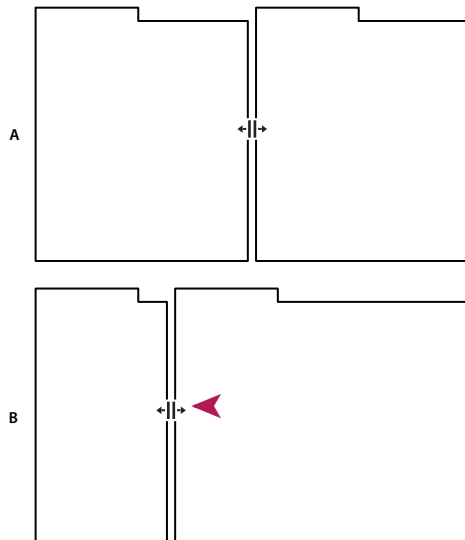
Dragging divider between panel groups to resize them horizontally
A. Original group with resize icon B. Resized groups

Open, close, and scroll to panels

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

- To open or close a panel, choose it from the Window menu.

- To close a panel or window, click its Close button ✖ .

- To see all the panel tabs in a narrow panel group, drag the horizontal scroll bar.

- To bring a panel to the front of a group of panels, do one of the following:

- Click the tab of the panel you want in front.

- Hover the cursor above the tab area, and turn the mouse scroll wheel. Scrolling brings each panel to the front, one after another.

- To reveal panels hidden in a narrow panel group, drag the scroll bar above the panel group.
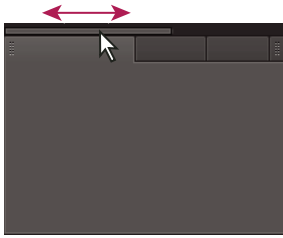
Drag horizontal scroll bar to see all panels in narrow group

Working with multiple monitors

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

More Help topics

"Dock, group, or float panels" on page 7

Position the Tools panel

The Tools panel provides quick access to tools, level meters, and the Workspace menu. By default, the Tools panel is a docked toolbar immediately below the menu bar. If you undock it, however, you can manipulate it like any other panel.

- To show or hide the Tools panel, choose Window > Tools.

- To undock the Tools panel from its default location, drag the handle at the left edge to another location in the workspace.

- To redock the Tools panel in its default location, drag the panel tab to the green drop zone that spans the entire width of the application window, just under the menu bar.

More Help topics

"Selecting audio" on page 33

"Level meters overview" on page 27

Display vertical and timeline rulers

In the Editor panel, the vertical ruler indicates amplitude in the waveform display and frequency in the spectral display. Timeline rulers indicate time location.

- To display the vertical ruler on the right side of the Editor panel, select View > Vertical Ruler.

- To add a second timeline ruler at the bottom of the Editor panel, choose View > Bottom Timeline Ruler.

If you show timeline rulers at top and bottom, each can display a different unit of time. (See "Change units in the timeline ruler or time display" on page 25.)

More Help topics

"About the waveform display" on page 28

"About the spectral display" on page 29

Change interface brightness and colors

1 Choose Edit > Preferences > Appearance (Windows) or Soundbooth > Preferences > Appearance (Mac OS).

2 Adjust any of the following options, and then click OK:

Brightness Brightens or darkens panels, windows, and dialog boxes.

Custom Colors Adjusts the color of waveforms, selections, and the current-time indicator.

Use Gradients When deselected, removes shadows and highlights from panels and buttons, and orange semicircles from beat markers.

More Help topics

"View layered or separated waveform channels" on page 30

"Position the current-time indicator" on page 25

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 Soundbooth

In the Window > Workspace submenu, Soundbooth provides the following predefined workspaces:

Default Optimizes the layout of panels for audio editing, providing a large view of the Editor panel.

Edit Audio to Video Places the Markers and Video panels above the Editor panel, so you can precisely synchronize audio and video.

Edit Score to Video Prominently positions the Video panel and increases the vertical size of the Properties panel, revealing all score options.

Meta Logging Extends the Metadata panel to the full height of the application window, helping you quickly edit properties for files, multitrack clips, and speech transcripts.

Save, reset, or delete workspaces

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, and 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 the current workspace to return to its original, saved layout of panels.

❖ 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: Importing, recording, and playing audio

Bring in audio from microphones, media files, and more. Then monitor it in stereo or surround-sound.

Configuring hardware inputs and outputs

Connecting to audio hardware

You can use a wide range of hardware inputs and outputs with Adobe Soundbooth. Sound card inputs let you bring in audio from sources such as microphones and tape decks. Sound card outputs let you monitor audio through devices such as speakers and headphones.
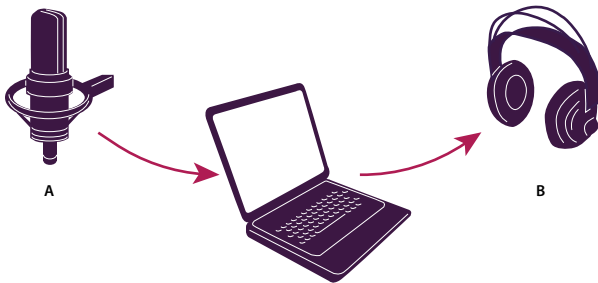
A. Sound card inputs connect to sources such as microphones and tape decks. B. Sound card outputs connect to speakers and headphones.

Set audio inputs and outputs

The audio inputs and outputs you select determine the default hardware ports Soundbooth uses for recording and playback. When recording, you can choose a different input device if necessary. (See "Record a new file" on page 17.)

1 Choose Edit > Preferences > Audio Hardware (Windows) or Adobe Soundbooth > Preferences > Audio Hardware (Mac OS).

2 For Default Device, choose a hardware interface. (For the best performance in Windows, choose an ASIO device. If none are available, choose Soundbooth 2.0 WDM or DirectSound.)

Note: In Mac OS, all possible combinations of inputs and outputs appear as separate devices. Choose System Default Input/Output unless you want Soundbooth to use different ports than other applications.

3 (Optional) Click Settings to set hardware driver properties. For more information, see one of the following:

• For a sound card in Mac OS, search for "Audio MIDI Setup" in Mac OS Help.

• For a professional ASIO sound card in Windows, see the documentation provided by the card manufacturer.

• For a WDM or DirectSound card in Windows, see "Set properties for standard Windows sound cards" on page 15.

4 In the Output Mapping section, specify the audio channel for each available hardware port.

Note: The Surround, Center, and LFE channels are available only if the default device is a multichannel audio interface.

5 (Mac OS only) For Buffer Size In Samples, optimize performance by choosing the lowest setting possible without audio dropouts. The ideal setting depends on the speed of your system, so some experimentation may be necessary.

6 If you want Soundbooth to continue playing audio when you switch to other applications, select Continue Audio Playback In Background.

Note: If multiple Windows applications use the same ASIO device, only one application can output audio at a time.

Set properties for standard Windows sound cards

In Windows, professional ASIO sound cards provide the best audio performance. If you're using a standard WDM or DirectSound card, however, you can improve performance by optimizing driver properties.

Note: The default driver properties work well for most systems. If you're experiencing slow response or audio dropouts, however, complete the appropriate procedure below.

Set properties for a WDM card

If you have a newer sound card that supports WDM, select this driver option. (WDM provides better performance than DirectSound.)

1 Choose Edit > Preferences > Audio Hardware.

2 From the Default Device menu, select Soundbooth 2.0 WDM Sound.

3 Click Settings.

4 On the Input and Output tabs, set the following options:

Enable Devices Determines which input and output ports are available in the Audio Hardware preferences.

Device 32-bit Recording and Playback Enables 32-bit inputs and outputs. To confirm whether a sound card supports this bit depth, see the documentation for the card.

Buffer Size Specifies the number of audio samples buffered during recording and playback. If you hear skips or dropouts, increase the buffer size; if playback and recording start slowly, decrease the size.

Set properties for a DirectSound card

If you have an older card that requires DirectSound, select this driver option.

1 Choose Edit > Preferences > Audio Hardware.

2 From the Default Device menu, select Soundbooth 2.0 DirectSound.

3 Click Settings.

4 In the DirectSound Full Duplex Setup control panel, set the following options, and then click OK.

Device check boxes Determine which ports are available in the Audio Hardware preferences.

Buffer Size (Samples) If you hear skips or dropouts, double-click the Buffer Size value for an input or output device, and type a higher value. If playback and recording start slowly, type a lower value.

Offset (Samples) Determines the audio latency (delay) between multiple sound cards. To enter a new value, double-click the current one.

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.

Sync Reference In systems with multiple sound cards, specifies which card sets the master clock. (The master clock synchronizes digital audio devices.)

Full Duplex Enables Soundbooth to simultaneously record and monitor audio. Leave this selected unless you have a very old sound card that doesn't support full-duplex operation.

Start Input First Determines the order in which Soundbooth starts the sound card input and output ports. Select this option only if you have a very old sound card that doesn't support full-duplex operation.

Opening, creating, and recording files

Open existing files

You can open audio and video in a variety of file formats, including AIFF, AVI, mp3, QuickTime, or WAV. If you open multiple files, the menu at the top of the Editor panel lets you choose which one to display.

Menu at top of Editor panel lets you choose which open file to display

1   Choose File > Open, or double-click the Files panel.

💡 *To add files to the Files panel without displaying them in the Editor panel, choose File > Import > Files.*

2   Select the desired files. (If you don't see the file you want, choose All Supported Media from the Files Of Type menu.)

3   Click Open.

💡 *For each open file, the Files panel displays properties such as Name, Media Type, and Duration. By default, files are sorted alphabetically by name. To sort files based on another property, click the property header.*

More Help topics

Browsing assets with Adobe Bridge

If you have the Adobe Creative Suite, Adobe Bridge helps you organize and browse assets for audio, video, and web projects. To access Adobe Bridge in Soundbooth, choose File > Browse In Bridge, or select a file in the Files panel and choose File > Reveal In Bridge.

With Adobe Bridge, you can do the following:

• View, search, and sort audio and video files

• Edit file metadata and keywords

• Rate and label files

Create empty audio files

To combine audio copied from multiple sources, create empty audio files.

1 Choose File > New > Empty Audio File.

2 Choose a sample rate, and select either Stereo or Mono.

More Help topics

"Copy, cut, and paste audio" on page 35

"Understanding sample rate" on page 4

Create files from selections and copied audio

Create a file from a selection

1 Make a selection in the Editor panel.

2 Choose File > New > Audio File From Selection.

Create a file from copied audio

1 Copy audio to the clipboard.

2 Choose File > New > Audio File From Clipboard.

More Help topics

"Create a multitrack file" on page 55

Record a new file

You can record audio from any device that you can plug in to your sound card's Microphone In or Line In port. Before recording, optimize your system's input levels, if you haven't already done so. (See "Adjust recording levels for standard sound cards" on page 18.)

1 Choose File > Record, or click the Record button ⬤ in the Editor panel.

2 Choose an audio input from the Device menu. If necessary, click Settings to set driver properties for the input.

3 Choose a Sample Rate option appropriate for your project. (See "Understanding sample rate" on page 4.)

4 Select Stereo or Mono, and then choose a Port option.

5 To ensure that the recorded signal is undistorted, select Monitor Input During Recording. (If you use this option for a microphone recording, wear headphones to avoid feedback.)

6 Enter a filename, and choose an incrementing scheme from the pop-up menu. (The incrementing scheme helps you distinguish between multiple related recordings with the same filename.)

7 To change the location for the saved file, click Browse.

8 Click the Record button ⬤ to begin recording.

9 As you record, click the Marker button ▽ to add audio markers (which you can later export as Adobe Flash cue points).

10 When you finish recording, click the Stop button ■, or close the dialog box.

Soundbooth automatically stores the new file in WAV format. If you prefer a different format, choose File > Save As.

More Help topics
"Use markers" on page 32

"Copying, pasting, cropping, and deleting audio" on page 35

Adjust recording levels for standard sound cards

Adjust levels if recordings are too quiet (causing background noise) or too loud (causing distortion). To get the best sounding results, record audio as loud as possible without clipping. When setting recording levels, watch the meters in the Record dialog box, and try to keep the loudest peaks in the yellow range below -3 dB.

Soundbooth doesn't directly control a sound card's recording levels. For a professional sound card, you adjust these levels with the mixer application provided with the card (see the card's documentation for instructions). For a standard sound card, you use the mixer provided by Windows or Mac OS.

Adjust sound card levels in Windows Vista and Windows 7

1 Right-click the speaker icon in the taskbar, and choose Recording Devices.

2 Double-click the input source you want to use.

3 Click the Levels tab, and adjust the slider as needed.

Adjust sound card levels in Windows XP

1 Double-click the speaker icon in the taskbar.

2 Choose Options > Properties.

3 Select Recording, and then click OK.

4 Select the input source you want to use, and adjust the Volume slider as needed.

Adjust sound card levels in Mac OS

1 Choose System Preferences from the Apple menu.

2 Click Sound, and then click the Input tab.

3 Select the device you want to use, and adjust the Input Volume slider as needed.

More Help topics
"Set audio inputs and outputs" on page 14

Configure and clean the media cache

To increase performance and import a wide variety of sample rates, Soundbooth creates cache files for each audio and video file you open. Soundbooth stores media cache files in a folder shared with other Adobe video applications. You can customize the location of this folder, clean the media cache database to improve performance, and delete cache files to conserve disk space.

1 Choose Edit > Preferences > Media (Windows) or Adobe Soundbooth > Preferences > Media (Mac OS).

2 Set the following options:

Media Cache Files To store cache files with source files rather than in the media cache folder, select Save Media Cache Files Next To Originals When Possible. To change the location of the media cache folder, click Browse.

Media Cache Database This MCDB file records the location of cache files and related media files. To remove references to files that are no longer in use, click Clean. To store the MCDB file in a different location, click Browse.

Delete Temporary .pek Files Removes small peak files that Soundbooth uses to quickly display audio waveforms.

Delete Temporary .cfa Files Removes large conformed audio files that Soundbooth creates when you open audio that uses a compressed format or an unsupported sample rate.

When Soundbooth is operating outside a file's original sample rate, parentheses surround the working sample rate in the Files panel. The Save command retains the original rate, if possible.

More Help topics

Ensure that Adobe video applications use the same cached files

Adobe video applications can automatically insert a unique document ID into each imported file. These unique IDs ensure that each application accesses the same cached previews and conformed audio files, preventing additional rendering and conforming.

❖ In the Media section of the Preferences dialog box (or the Metadata section for Soundbooth), select Write XMP IDs To Files On Import.

This setting is global—a change in one Adobe video application affects all the others. This setting also results in new file modification dates when IDs are initially inserted.

To save rendering time when transferring a project to another computer, move both cached and original files.

Viewing and editing XMP metadata

About the Metadata panel and XMP metadata

To streamline your workflow and organize your files, use *XMP metadata*. Metadata is a set of descriptive information about a file. Video and audio files automatically include basic metadata properties, such as date, duration, and file type. You can add details with properties such as location, director, copyright, and much more.

With the Metadata panel, you can share this information about assets throughout Adobe video and audio applications. Unlike conventional clip properties, which are limited to only one application's Project or Files panel, metadata properties are embedded in source files, so the data automatically appears in other applications. This sharing of metadata lets you quickly track and manage video assets as they move through your production workflow.

Note: Properties in the Metadata panel also appear in Adobe Bridge, providing additional details that help you quickly browse assets.

For a video about the Metadata panel, see www.adobe.com/go/lrvid4104_xp.

About schemas and properties

A metadata *schema* is a collection of properties specific to a given workflow. The Dynamic Media schema, for example, includes properties like Scene and Shot Location that are ideal for digital video projects. Exif schemas, by contrast, include properties tailored to digital photography, like Exposure Time and Aperture Value. More general properties, like Date and Title, appear in the Dublin Core schema. To display different properties, see "Show or hide XMP metadata" on page 21.

🔆 *For information about a specific schema and property, hover the pointer over it in the Metadata panel. For most items, a tool tip appears with details.*

About the XMP standard

Adobe applications store metadata using the Extensible Metadata Platform (XMP). XMP is built on XML, which facilitates the exchange of metadata across a variety of applications and publishing workflows. Metadata in most other formats (such as Exif, GPS, and TIFF) automatically transfers to XMP so you can more easily view and manage it.

In most cases, XMP metadata is stored directly in source files. If a particular file format doesn't support XMP, however, metadata is stored in a separate *sidecar file*.

Project assets without corresponding files don't support XMP. Examples from Adobe Premiere Pro include Bars and Tone, Universal Counting Leader, Color Matte, Titles, Black Video, and Transparent Video.

🔆 *To customize the creation and exchange of metadata, use the XMP Software Development Kit. For more information about XMP, see Extensible Metadata Platform.*

About file, clip, and project XMP metadata

For the most part, Adobe video and audio applications deal with XMP metadata very similarly. Some small distinctions exist, however, reflecting the unique workflow stage that each application addresses. When using applications in tandem, an understanding of these slightly different approaches can help you get the most out of metadata.

Adobe OnLocation and Encore provide one set of metadata properties for all assets. However, Adobe Premiere Pro, After Effects, and Soundbooth divide the Metadata panel into separate sections for different asset types.

Adobe Premiere Pro Separates metadata in these sections:

• **Clip** Displays properties for clip instances you select in the Project panel or Timeline panel. This metadata is stored in project files, so it appears only in Adobe Premiere Pro.

• **File** Displays properties for source files you select in the Project panel. This metadata is stored directly in the source files, so it appears in other applications, including Adobe Bridge.

After Effects Separates metadata in these sections:

• **Project** Displays properties for the overall project. If you select Include Source XMP Metadata in the Output Module Settings dialog box, this information is embedded into files you output from the Render Queue.

• **Files** Displays properties for source files you select in the Project panel. (If you select a proxy, properties for the actual file appear.)

For After Effects, both Project and File properties are stored directly in files, so you can access this metadata in Adobe Bridge.

Soundbooth Separates metadata in these sections:

• **File** Displays properties for the currently displayed audio or ASND file. This metadata is stored directly in such files, so it appears in other applications. (Adobe Bridge, however, does not display metadata for ASND files.)

- **Clip** Displays properties for multitrack clips you select in the Editor panel. This metadata is stored in the containing ASND file, so it appears only in Soundbooth.

Adobe Premiere Pro and Soundbooth also provide a Speech Analysis section with metadata that appears only in those applications.

Show or hide XMP metadata

To optimize the Metadata panel for your workflow, show or hide entire schemas or individual properties, displaying only those that you need.

1 From the options menu ▼≡ for the Metadata panel, select Metadata Display.

2 To show or hide schemas or properties, select or deselect them from the list.

Save, switch, or delete metadata sets

If you use multiple workflows, each requiring different sets of displayed metadata, you can save sets and switch between them.

1 From the options menu ▼≡ for the Metadata panel, select Metadata Display.

2 Do any of the following:

- To save a customized set of displayed metadata, click Save Settings. Then enter a name, and click OK.

- To display a previously saved set of metadata, select it from the menu.

- To delete a previously saved set of metadata, select it from the menu, and click Delete Settings.

Create schemas and properties

If you have a unique, customized workflow that the default metadata options don't address, create your own schemas and properties.

1 From the options menu ▼≡ for the Metadata panel, select Metadata Display.

2 Click New Schema, and enter a name.

3 In the list, click Add Property to the right of the schema name.

4 Enter a property name, and select one of the following for Type:

Integer numbers that you drag or click to change

Real fractional numbers that you drag or click to change

Text text box (for properties similar to Location)

Boolean check box (for On or Off properties)

Search XMP metadata

1 Select the files or clips you want to search.

2 In the search box at the top of the Metadata panel, enter the text you want to find.

The list of metadata collapses to reveal only properties that contain your search string.

3 (Adobe Premiere Pro only) To navigate through the search results, click the Previous and Next buttons ◁ ▷ to the right of the search box, or press Tab.

4 To exit the search mode and return to the full list of metadata, click the close button ✕ to the right of the search box.
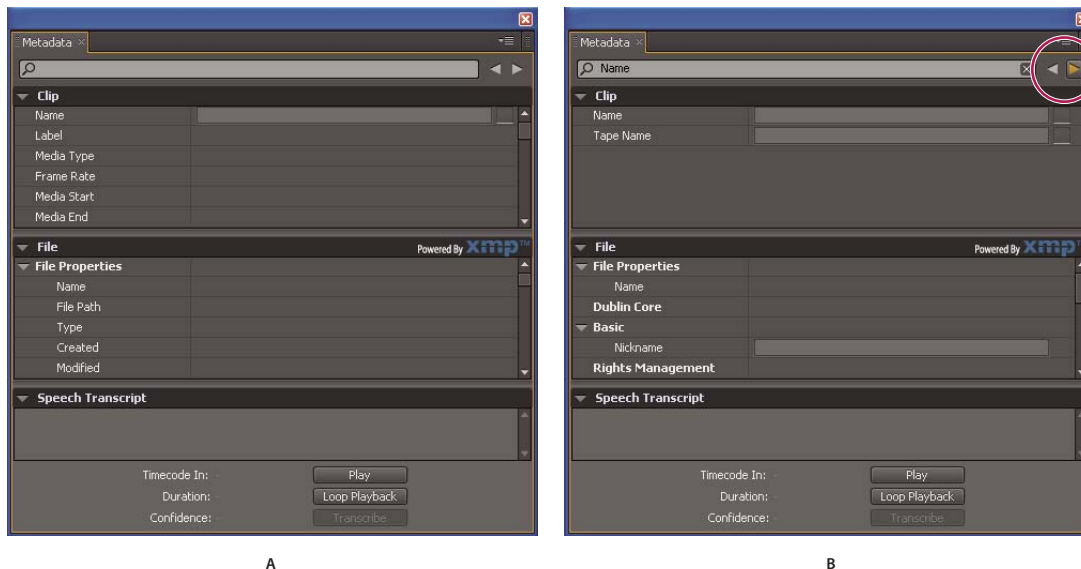
Metadata panel
A. Before search, all properties appear   B. After search, only properties with search string appear. In Adobe Premiere Pro, Previous and Next buttons navigate through search results.

Edit XMP metadata

In Adobe video applications, similarly named properties are linked in the Metadata and Project panels. However, the Metadata panel provides more extensive properties and lets you edit them for multiple files simultaneously.

Note: Instead of a Project panel, Soundbooth uses the Files panel.

1 Select the desired files or clips.

2 In the Metadata panel, edit text or adjust values as needed.

  If you selected multiple items, the panel displays properties as follows:

  • If a property matches for all items, the matching entry appears.

  • If a property differs, *<Multiple Values>* appears. To apply matching values, click the text box, and type.

Analyze speech for text XMP metadata

Adobe Premiere Pro and Soundbooth analyze spoken words and generate text metadata. You can edit and search text metadata like any other metadata properties. You can then navigate to the times at which specific words are spoken, to better align edits, advertising, and subtitles.

For more information, see the video tutorial Using Speech Search to Speed Editing.

Note: Useful results from speech analysis require good audio quality. Background noise significantly reduces accuracy. To remove such noise, use the tools and processes in Soundbooth.

More Help topics

Speech to text enhancements

Analyze speech to create text metadata

1 Select a file or clip.

2 At the bottom of the Metadata panel, click Analyze Speech, or Analyze (Adobe Premiere Pro).

3 Set the Language and Quality options, and select Identify Speakers if you want to create separate speech metadata for each person.

Note: Speech Search can use any of several language-specific and dialect-specific libraries, such as libraries for Spanish and UK English.

4 Click OK.

The spoken words appear in the Speech Analysis section.

5 To retain the speech metadata, save the project.

If you import files with a speech metadata into After Effects, each word appears as a layer marker on layers based on these footage items.

Navigate to a specific word in speech metadata

1 In the Speech Analysis section, select the word.

Timecode In and Duration indicate the precise location and length of your selection.

2 To hear the selection, click either Play or Loop Playback. (The latter option repeatedly plays the selected word, with some preroll and postroll.)

Edit speech metadata

❖ In the Speech Analysis section, do any of the following:

- To correct a word, click it, and type.

- To insert, delete, merge, cut, or copy words, right-click an existing word, and choose a command from the context menu.

Copy text from speech metadata to the clipboard for use in a text editor

❖ Right-click the transcript, and choose Copy All.

Improve speech analysis with reference scripts

Accuracy of the speech to text conversion depends on the clarity of the spoken words and the quality of the recorded dialog. Dialog recorded in a noisy environment or with poor microphone placement cannot produce highly accurate results even with a reference script. You can nevertheless use a reference script to improve speech analysis. A reference script is a text document containing dialog similar to the dialog recorded in your assets.

There are two types of reference scripts:

- A script that contains similar dialog, but was not necessarily written for the current project. For example, a series of medical training scripts for different products can be combined into a single text document. You can use this text document as a reference script. With this type of reference script, speech analysis produces results more accurate than it does when using only the default language models.

- A script that matches the recorded dialog. This type of reference script provides the highest accuracy possible. For example, you can use the script that the talent read during the shoot as a reference script. Alternatively, you could use a transcript typed from the assets for the purposes of close captioning.

Speech Analysis supports reference scripts only in the UTF-8 encoded text format, including Adobe Story scripts, which have the .astx filename extension.

Note: To make reference scripts available in Soundbooth, first complete the steps below in Adobe Premiere Pro.

1 From the Reference Script menu in the Analyze Content dialog box, choose Add.

2 In the displayed dialog box, browse to the reference script text or .astx file, select it, and click OK.

3 In the Import Script dialog box, type a name for the reference script, and select the language of the script.

   Note: You can view the text of the file in a scrolling window.

4 Select Script Text Matches Dialog only if the imported script covers the recorded dialog verbatim. For example, if the reference script is the script from which the talent read their lines, select Script Text Matches Dialog. Select this option even if the recorded dialog is shorter than what the script file covers.

5 Click OK.

   The Import Script dialog closes, and the reference script is selected in the Reference Script menu.

6 Click OK.

Improve speech analysis with Adobe Story, OnLocation, and Adobe Premiere Pro

You can use Adobe Story, OnLocation, and Adobe Premiere Pro to create the most accurate speech analysis. Import a script written in Adobe Story into OnLocation. OnLocation produces a list of shot placeholders for each scene. Either record these shots using OnLocation during production, or link the placeholder shots to their respective video files when you import the video files into OnLocation. In either case, OnLocation embeds the text for each shot from the original script into the metadata of the shot.

When you import the clips into Adobe Premiere Pro, it automatically uses the Adobe Story script as a reference script. When Adobe Premiere Pro finds enough matches with the embedded script, Adobe Premiere Pro replaces the analyzed speech text with the embedded script text. Adobe Premiere Pro carries over correct spelling, proper names, and punctuation from the reference script, benefits that standard speech analysis cannot provide.

The closeness of the match between the embedded script text and the recorded dialog determines the accuracy of matched-script text. If 100% accuracy is important, edit and revise the script text first. Ensure that the script matches the recorded dialog before using it as a reference script.

Playing audio

Monitoring time during playback

The Editor panel provides several features to help you monitor time during playback:

- The timeline ruler displays hours, minutes, and seconds by default.
- The current-time indicator 🛡 determines the starting point for playback and moves through the waveform as you listen to audio files.
- The time display shows the current time in numerical format.
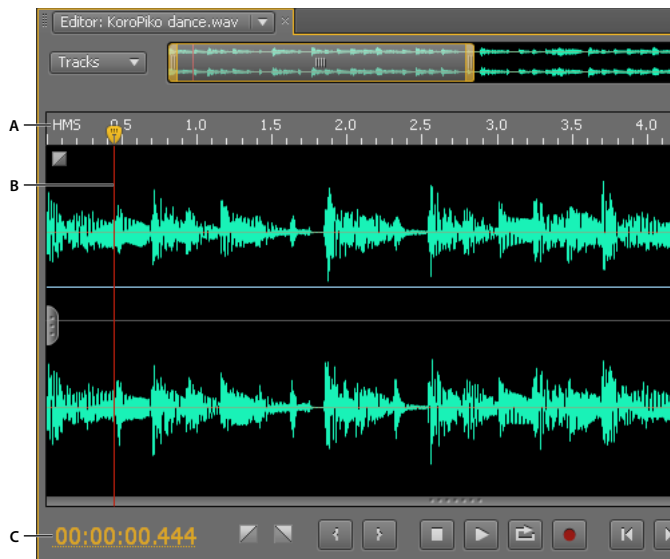
Features that help you monitor time
A. Timeline ruler   B. Current-time indicator   C. Time display

More Help topics

"Display vertical and timeline rulers" on page 11

Change units in the timeline ruler or time display

❖ Right-click the timeline ruler or time display, and select one of the following:

HMS (Timeline ruler only) Shows hours, minutes, and seconds.

Decimal (HH:MM:SS:mmm) Shows hours, minutes, seconds, and milliseconds.

Samples (HH:MM:SS:sample) Shows hours, minutes, seconds, and audio samples.

Samples Shows only audio samples. (This option is unavailable for multitrack files.)

To better understand samples, see "Digitizing audio" on page 4.

FPS options Show frames-per-second for various film and video formats, including high-definition options ranging from 50–60 fps. DF and NDF indicate drop-frame and non-drop frame formats, respectively.

Feet + Frames options (Time display only) Show location in traditional 16mm or 35mm film. You should display these units when editing a digital proxy in a film-based project.

Custom Shows the custom time format specified in the Preferences dialog box. The default, 12 frames-per-second, matches the default in Adobe Flash Professional. If your Flash projects use a different frame rate, choose Edit Custom Time Format.

Position the current-time indicator

• In the timeline ruler, click to instantly reposition the current-time indicator 👤. Or, drag the indicator to *scrub* audio, previewing it at different time points.

• In the time display, drag the numbers, or click them to enter a specific time.

- In waveform or spectral display, click or drag with a selection tool. (The current-time indicator is placed at the beginning of the selection.)

💡 *To hear a selection in the context of surrounding audio, disable the Loop Playback option in the transport controls, and position the current-time indicator prior to the selection.*
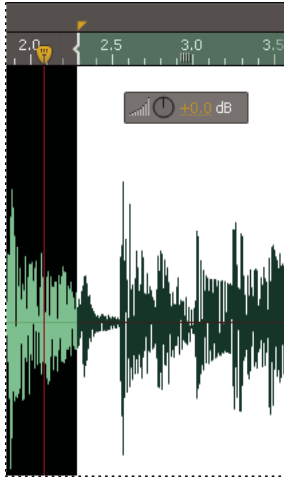
Position the current-time indicator prior to selection to hear surrounding audio.

More Help topics

"Change interface brightness and colors" on page 12

Transport controls

At the bottom of the Editor panel, Soundbooth provides several transport controls for adjusting playback.

💡 *To start or stop playback without using the transport controls, press the spacebar.*

Go To Previous ◀| Moves the current-time indicator to the previous marker, the previous selection edge, or the beginning of the file.

Go To Next |▶ Moves the current-time indicator to the next marker, the next selection edge, or the end of the file.

Stop ■ Stops playback.

Play ▶ Starts playback at the current-time indicator. By default, this button becomes the Pause button ❚❚ after playback begins. To return the current-time indicator back to its original position, click the Stop button.

💡 *To reflect the playback behavior of Adobe Premiere Pro, deselect Return To Start Position On Stop in the General section of the Preferences dialog box.*

Loop 🔁 Enables looped playback of audio.

Record 🔴 Opens the Record dialog box. (See "Record a new file" on page 17.)

More Help topics

"Shortcuts for playing and zooming audio" on page 75

Level meters overview

To monitor audio amplitude, use the level meters. When you edit individual files, meters appear only in the toolbar. When you mix multitrack files, meters in the toolbar display overall levels, while meters in the Editor panel display track levels.

For stereo files, the upper meter represents the left channel, and the lower meter represents the right channel. For surround-sound (5.1) files, six meters represent the following channels from top to bottom: Left, Right, Left Surround, Right Surround, Center, and LFE.

Level meters for stereo file
A. *Left channel* **B.** *Right channel* **C.** *Peak indicator* **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. All levels below that maximum are expressed as negative numbers.

To identify amplitude at a specific point in the meters, position the mouse over that point, and note the value in the tool tip. To the right of the meters, Soundbooth indicates the peak level since playback last began. If clipping occurs, this value is replaced by the word *Over*, and red clip indicators appear in the meters.

To clear clip indicators and peak amplitude values, click them, or right-click the meters and choose Reset Meters.

More Help topics

"Understanding bit depth" on page 5

Chapter 5: Editing and repairing audio files

The intuitive visual tools in Adobe Soundbooth make it easy to edit, optimize, and repair audio.

Displaying audio

View audio waveforms and spectrums

When you open an audio file, the Editor panel provides a 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 Editor panel.

Below the panel's default waveform display, which is ideal for evaluating audio amplitude, you can view audio in the spectral display, which reveals audio frequency (low bass to high treble).

❖ To view the spectral display, do any of the following:

• In the toolbar, click the Spectral Frequency Display button [image].

• In the Tasks panel, click Remove A Sound.

• In the Editor panel, drag the divider between the waveform and spectral displays to gradually change the proportion of each. To instantly show or hide the spectral display, double-click the handle or click the triangle to its right.
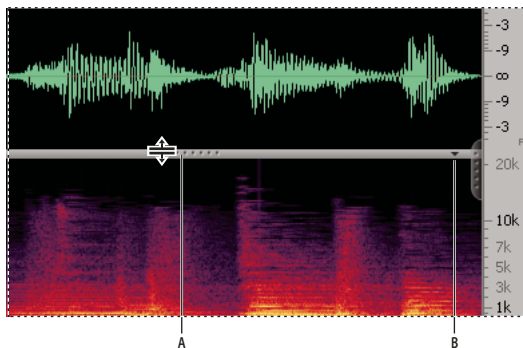
Viewing the waveform and spectral displays
A. *Drag the divider to change the proportion of each.* ***B.*** *Click the triangle to show or hide the spectral display.*

More Help topics

"Waveform measurements" on page 2

About the waveform display

The waveform display shows audio as a series of amplitude peaks and valleys. The x-axis (horizontal ruler) measures time, and the y-axis (vertical ruler) measures amplitude on a decibel scale that ranges from −∞ (negative infinity) for silence to 0 dBFS for loud peaks. Quiet audio has both lower peaks and lower valleys than loud audio.

With its clear indication of amplitude changes, the 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 syllable.
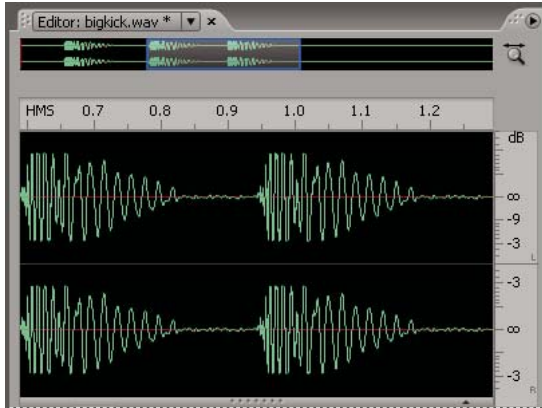

Stereo file in waveform display

More Help topics

About the spectral display

The spectral display shows audio 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. Colors represent amplitude, ranging from dark blue for low amplitude to bright yellow for high amplitude.

The spectral display is perfect for removing unwanted sounds, such as clicks, coughs, buzz, hum and other artifacts. This is known as *frequency-space editing*.
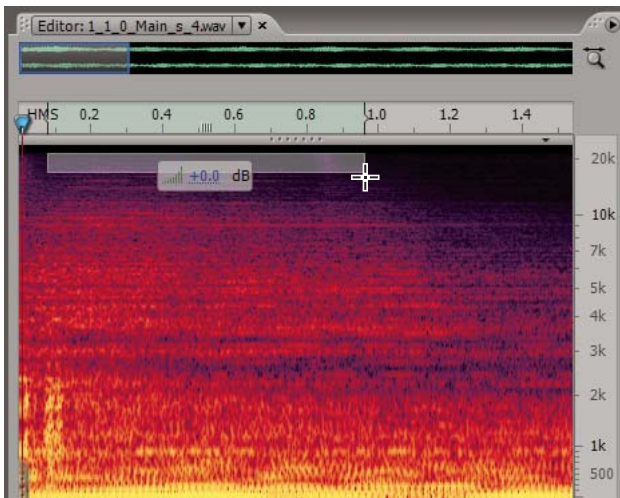
Spectral display, with high frequencies selected
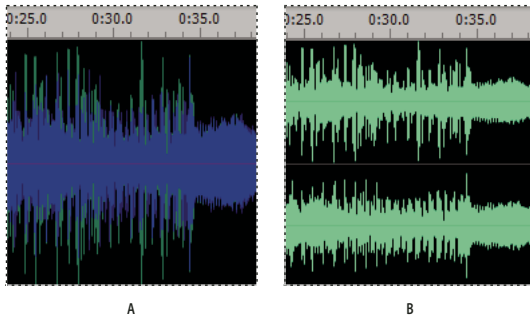
More Help topics

View layered or separated waveform channels

For stereo and surround-sound files, you can view layered or separated channels. Layered channels overlay each other in different colors, better revealing overall volume changes. By contrast, separated channels appear individually, better revealing distinct volume changes.

❖ Choose View > Channels, and then choose Layered or Separated.

Channel View options
A. *Layered* **B.** *Separated*

More Help topics

"Change interface brightness and colors" on page 12

Customize the spectral display

1 In the Tasks panel, click Remove A Sound.

2 From the Resolution menu, choose a high setting to display frequencies more accurately, but time location less accurately. Or, choose a low setting to do the opposite.

💡 *Use high Resolution settings to identify artifacts of long duration (like squeaks or 60-Hz hum). Use low settings to identify transient peaks (like clicks and pops).*

3 For Vertical Scale, enter a higher number to display frequencies more logarithmically, or a lower number to display them more linearly.

💡 *A logarithmic display better reflects the uneven frequency emphasis of human hearing, providing better low-frequency detail. A linear display gives equal weight to each frequency.*

4 To limit playback to frequencies you've selected with the Frequency Selection, Rectangular Marquee, or Lasso tools, select Play Selected Frequencies Only.

By default, the option above is deselected, so Soundbooth plays all frequencies in the same time range as selected frequencies.

More Help topics

"Visually identifying noise" on page 39

Zoom audio

Zoom into the current display

Place the pointer over the Editor panel, and roll the mouse wheel. (Roll over the waveform display to zoom time ranges; roll over the spectral display to change the vertical scale.)

Zoom into a selected area

In the toolbar, select the Zoom tool 🔍 , then click or drag in the Editor panel. (To zoom out, Alt-click or Option-click.)

Zoom into a specific time range

In either the zoom navigator or the timeline ruler, right-click and drag. The magnifying glass icon ⇄🔍 creates a selection showing the range that will fill the Editor panel.

Zoom into a specific frequency range

In the vertical ruler for the spectral display, right-click and drag.
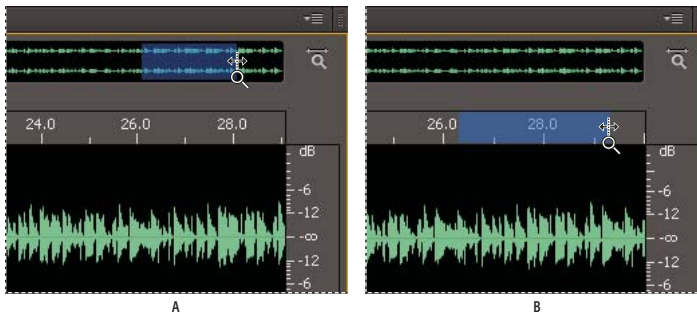
Extend or shorten the displayed range

Place the pointer over the left or right edge of the highlighted area in the zoom navigator, and then drag the magnifying glass icon.

Magnify selected audio

From the View menu, choose Zoom In At In Point, Zoom In At Out Point, or Zoom To Selection.

Display the entire audio file

In the Editor panel, click the Zoom Out Full button 🔍 .

To zoom into a specific time range, right-click and drag
A. Zoom navigator B. Timeline ruler

More Help topics

"Shortcuts for playing and zooming audio" on page 75

Navigate through time

At higher zoom levels, you may need to navigate to different audio content in the Editor panel.

• Select the Hand tool ✋ in the toolbar, then drag in the Editor panel.

• In the zoom navigator, drag left or right.

Scrolling with the zoom navigator

More Help topics

"Position the current-time indicator" on page 25

Use markers

Markers make it easy to navigate in a waveform, perform edits, or play back audio. A marker refers to a specific time position (for example, 00:08:07.566 from the start of a file).

Though you can add and move markers in the Editor panel, the Markers panel provides additional options. To show or hide the Markers panel, choose Window > Markers.
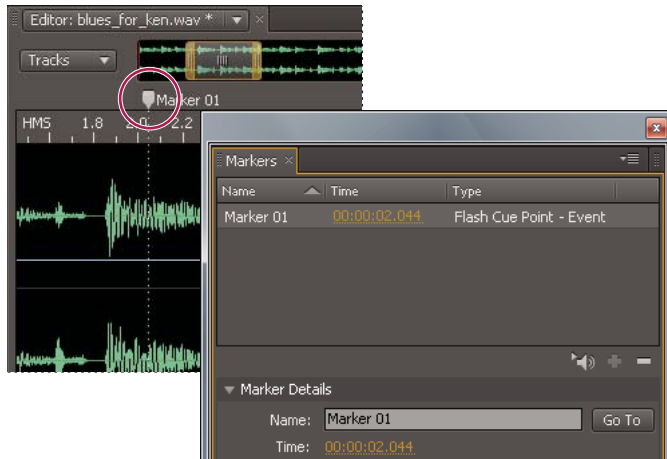
Note: To add and adjust markers in multitrack clips, use single-clip editing mode. (See "Edit and process a single multitrack clip" on page 61.)

A marker in the Markers and Editor panels

Add markers

1 Either start playback to add general markers, or position the current-time indicator 🔱 to add a marker in a specific location.

2 Do any of the following:

   • Press the asterisk (*) key on the numeric keypad or the *M* key.

   • Choose Edit > Marker > Set Flash Cue Point.

   • Click the Add Marker button ➕ in the Markers panel

💡 *To add markers while recording audio, see "Record a new file" on page 17.*

Move a marker

• In the Editor panel, drag the marker to a new location.

• In the Markers panel, enter a new Time value.

Go to a marker in the waveform

❖ In the Markers panel, double-click the marker.

💡 *To automatically play markers when you go to them, click the Auto Play button 🔊. This option helps you quickly identify marker locations.*

Rename a marker

❖ In the Markers panel, select the marker, expand the Marker Details section, and enter a new name in the Name text field.

Delete markers

• To remove specific markers, select them in the Markers panel, and click the Clear Marker button ▭.

• To remove all markers, choose Edit > Marker > Clear All Markers.

More Help topics

"Set properties for cue points" on page 65

"Export or import cue points in XML files" on page 66

Selecting audio

Select time ranges

1 In the toolbar, select the Time Selection tool ⌶.

2 In the Editor panel, drag to select a specific range, or double-click to select the visible range.

💡 *To select a precise timecode range, position the current-time indicator 🛡, and click the Set In Point ▣ and Set Out Point ▣ buttons.*

3 (Optional) Do any of the following:

• To extend or shorten a selection, drag the left or right edge of the highlighted range. Or, drag the left or right selection edge in the timeline ruler.

• To move the selection without changing its length, drag the middle of the selection in the timeline ruler.
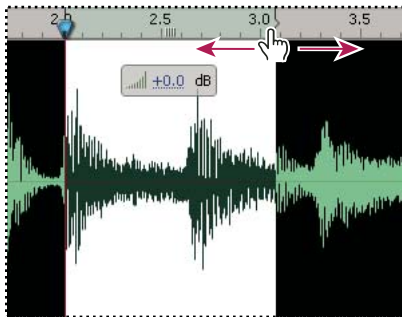
Extending a selection in the timeline ruler

💡 *To edit stereo channels separately, choose File > New > Multitrack File From Channels.*

More Help topics

"Position the current-time indicator" on page 25

"Create files from selections and copied audio" on page 17

"Change interface brightness and colors" on page 12

Select frequency ranges

In the spectral display, three tools let you select audio data within specific frequencies. The Frequency Selection tool selects an entire frequency range throughout a file, the Marquee tool selects a rectangular area, and the Lasso tool creates free-form selections. All three tools provide powerful flexibility for audio restoration projects. For example, if you find a sonic imperfection, you can select and edit just the affected frequencies, with superior results and faster processing.

1   In the toolbar, click the Spectral Frequency Display button ▪▪▪. Then select the Frequency Selection ⊱,
    Marquee ⸢⸥, or Lasso ◯ tool.

2   In the spectral display of the Editor panel, drag to select audio.
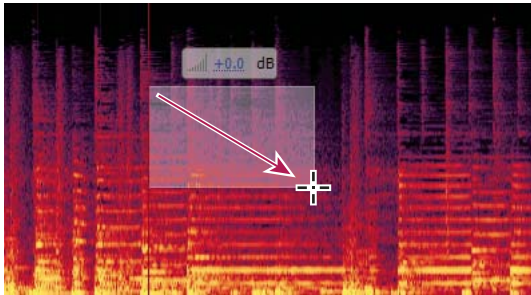
Marquee selection in the spectral display

3   (Optional) Do any of the following:

•   To extend or shorten a selection, drag its left or right edge in the timeline ruler.

•   To move the selection without changing its length, drag the middle of the selection in the timeline ruler.

More Help topics

"About the spectral display" on page 29

"Repairing audio" on page 39

"Create files from selections and copied audio" on page 17

Select all of a waveform

•   To select the visible range of a waveform, double-click in the Editor panel. Or choose Edit > Select View.

•   To select an entire file, triple-click in the Editor panel. Or choose Edit > Select All.

If nothing is selected, Soundbooth applies effects and processes to entire files.

More Help topics

"Save entire files or selected ranges" on page 69

Snap to markers, beats, frames, or the timeline ruler

To more accurately position selections and the current-time indicator, enable snapping.

❖   Choose View > Snapping, and then select Snap to Ruler, Markers, Beats, or Frames.

More Help topics

Copying, pasting, cropping, and deleting audio

Copy, cut, and paste audio

1 With the Time Selection tool I , select the audio you want to copy or cut.

2 To copy audio to the clipboard, choose Edit > Copy. Or, to remove audio from the current file and add it to the clipboard, choose Edit > Cut.

💡 *When you edit audio from a video file, the Cut command silences audio without changing file length, maintaining synchronization with video.*

3 In any file, place the current-time indicator 🎯 where you want to insert audio, or select the audio you want to replace. Then choose Edit > Paste.

More Help topics

Mix audio when pasting

The Mix Paste command lets you mix audio from the clipboard with the current waveform.

1 In the Editor panel, place the current-time indicator 🎯 where you want to start mixing audio. Or, select the audio you want to replace.

2 Choose Edit > Mix Paste.

3 Drag the sliders for Copied Audio and Existing Audio to adjust the ratio of each.

4 Click Preview, and adjust the sliders as needed.

5 Click OK to apply your changes.

More Help topics

Trim, crop, or delete audio

Trim audio at the beginning or end of a file

Click the Zoom Out Full button 🔍 to display trim handles in the Editor panel, and drag either handle inward.

💡 *If trim handles are visually distracting, deselect View > Trim Handles.*

Crop to selected audio

Select the audio you want to keep, and then choose Edit > Crop.

Delete audio in the middle of a file

Select the audio you want to remove, and then press the Delete key.

💡 *When you edit audio from a video file, trimming, cropping, and deleting silence audio without changing file length, maintaining video synchronization. (See "Maintaining video synchronization when editing audio" on page 67.)*
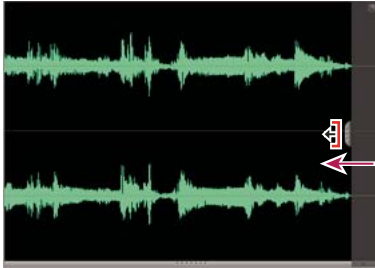
Use trim handles to quickly remove the beginning or end of files.

More Help topics

"Selecting audio" on page 33

"Position the current-time indicator" on page 25

Fading, changing, and maximizing volume

Automatically fade in or out

1 Do one of the following:

• To fade to a specific point, select a range from that point to either end of the file.

• To apply a five-second fade, move the current-time indicator 🛡 to either end of the file.

2 At the bottom of the Editor panel, click the Fade In button ◣ or the Fade Out button �◣.

More Help topics

"Select time ranges" on page 33

"Trim, crop, or delete audio" on page 35

"Insert silence" on page 39

Precisely fade in or out

1 At the far left or right of the waveform, drag the Fade In ◣ or Fade Out ◣ handle inward.

2 (Optional) To change from the default, linear fade, drag up to create a quick, smooth logarithmic fade, or drag down to create a long, smooth exponential fade.

💡 *If you save files in ASND format, fades remain adjustable and appear in multitrack clips.*
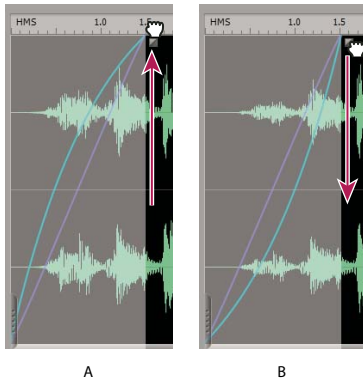
Drag up or down to change the fade type
A. *Quick, logarithmic fade* **B.** *Long, exponential fade*

More Help topics

"Trim, crop, or delete audio" on page 35

"Insert silence" on page 39

Raise or lower volume

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

2 Drag the numbers in the volume control that appears above selected audio.

Note: The 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.
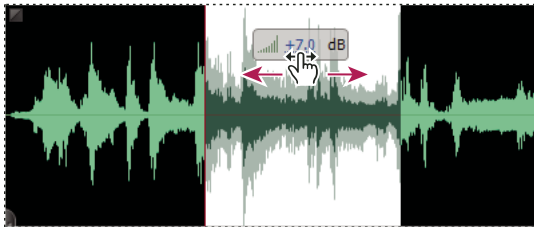
Changing the volume of a selected area

To change volume or pan audio in multitrack mixes, see "Mix tracks and clips" on page 58.

More Help topics

"Automate mixes with keyframes" on page 59

"Compressor effect" on page 49

"Dynamics effect" on page 51

Maximize volume

After you edit audio and apply effects, maximize volume as a finishing touch. Soundbooth offers two techniques that raise volume to 0.3 dBFS, just below the digital maximum, ensuring optimal volume while avoiding clipping. *Normalizing* retains dynamic range by amplifying an entire file equally. *Hard limiting* reduces dynamic range by amplifying quieter sounds more than loud ones.

Apply hard limiting to increased perceived volume and make your audio stand out in web, video, or radio presentations.

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

2   At the bottom of the panel, click the Louder button 🔊 once to normalize audio, or multiple times to apply hard limiting.

With each successive click of the button, Soundbooth increases overall volume by 3 dB. Hard limiting prevents clipping.

To separately apply Normalize and Hard Limiting commands, use the Processes menu.

More Help topics

"dBFS" on page 78

"Compressor effect" on page 49

"Dynamics effect" on page 51

Equalize volume within a file

To make an entire selection or file sound equally loud, do the following:

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

2   At the bottom of the panel, click the Equalize Volume Levels button 🔊.

Match volume across multiple files

If you plan to present a group of audio files on CD, the web, or elsewhere, use the Match Volume task to make them sound consistent.

1   Choose Tasks > Volume Correction.

2   In the Tasks panel, click the Match Volume tab.

3   Either drag files to the Files To Match list, or click the Add Files icon 🔲 below the list and navigate to files.

4   (Optional) To display average, perceived, and peak volume for each file, click the Compute icon 🔍. These statistics help you specify target values in the next step.

5   Select one of the following options:

Match To Volume Matches to an average value you specify. Choose this option if no file in the group reflects the desired volume.

Match To File Matches to the average volume of a specific file. Choose this option if one file reflects the target volume you want to achieve.

The Account For Perceived Loudness option focuses analysis on midrange frequencies that the ear is most sensitive to. Keep this option selected unless frequency emphasis varies greatly in a file (for example, midrange frequencies are pronounced in a short passage, but bass frequencies are elsewhere).

Match To Peak Volume Matches to a peak value you specify. Because this option retains dynamic range, it's a good choice for files you plan to process further, or for highly dynamic audio like classical music.

6 Click Match Volume.

7 Review the reported volume change in each file, and click Close.

Insert silence

Insert silence to separate different types of program material, or to synchronize an audio file with live broadcasts or video.

1 Do either of the following:

- To silence existing audio, select it.

- To add a specific duration of silence, position the current-time indicator 🛡 where you want silence to begin.

2 Choose Edit > Insert > Silence. If you're adding silence, specify a duration in seconds, and click OK.

More Help topics

"Trim, crop, or delete audio" on page 35

"Raise or lower volume" on page 37

Repairing audio

Visually identifying noise

In the spectral display, you can quickly identify and select different types of noise. Clicks and crackle, for example, usually appear as bright vertical bars that extend from top to bottom. Hiss, meanwhile, appears as a light-red cloud that extends across the top.
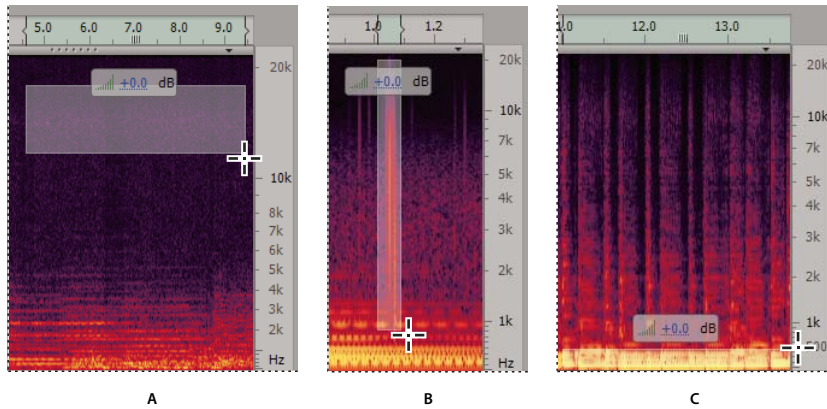
Selecting various types of noise in the spectral display
A. Hiss B. Crackle C. Rumble

For a video about repairing audio, see *www.adobe.com/go/vid0245*.

More Help topics

"About the spectral display" on page 29

"Select frequency ranges" on page 34

Remove background noise

To remove steady background noise such as hiss or hum, use the Noise process.

1   (Optional) Help Soundbooth better distinguish noise from desirable audio. In the Editor panel, select an area that contains only noise. Then choose Processes > Capture Noise Print.

2   In the Editor panel, select the audio you want to clean up.

3   Choose Processes > Reduce Noise.

4   Set the following options:

Reduction Specifies the amplitude of the *noise floor*. Audio that falls below this threshold is attenuated. Use lower values for audio with a wide dynamic range, such as classical music with loud and quiet passages. Use higher values for audio with a compressed dynamic range, such as pop music.

Reduce By Determines how much to attenuate audio that falls below the noise floor.

Use Captured Noise Print Distinguishes noise from desirable audio using the noise print you captured in step 1.

5   Click Preview, and adjust settings as needed. To compare processed and original audio, click the power button ⏻ .

6   When you achieve the best-sounding results, click OK to process the audio file.

More Help topics

"Selecting audio" on page 33

Remove clicks, pops, or rumble

Soundbooth quickly removes sharp clicks and pops (such as crackle from vinyl records), or low-end rumble (such as vibrations from passing trucks).

1 In the Editor panel, select the audio you want to clean up.

2 From the Processes menu, choose Remove Clicks & Pops or Remove Rumble.

3 Drag the appropriate slider below:

Click/Pop Determines sensitivity to audio artifacts. Higher settings detect more artifacts but may also remove audio you wish to retain. Use lower values to remove subtle clicks or high values to remove loud pops.

Rumble Determines the cut-off frequency for rumble, removing all lower frequencies. Try lower values for deep rumble such as traffic noise, and higher values for rumble such as microphone handling sounds.

4 Click Preview, and adjust settings as needed. To compare processed and original audio, click the power button ⏻.

5 When you achieve the best-sounding results, click OK to process the audio file.

More Help topics

"Selecting audio" on page 33

Remove individual sounds

Unlike the Delete command, which can introduce audible glitches, the Auto Heal command removes an unwanted sound and seamlessly blends the area with surrounding audio.

💡 *To automatically heal the outer edges of all edited selections, select Auto-Heal Edit Boundaries in the General section of the Preferences dialog box.*

1 Choose Tasks > Remove A Sound.

2 In the Tasks panel, select the Time Selection ⌶ , Frequency Selection ⊶, Marquee ⬚ or Lasso ⬭ tool.

3 In the Editor panel, zoom in on the sound you want to remove.

4 In the spectral display, make a selection no longer than 25,000 samples (.52 seconds at a sample rate of 48 kHz).

💡 *To precisely determine selection length, right-click the timeline ruler, and choose Samples.*

5 In the Tasks panel, click Auto Heal. Or choose Processes > Auto Heal.

More Help topics

"Zoom audio" on page 30

"Selecting audio" on page 33

"Customize the spectral display" on page 30

Looping, stretching, and pitch shifting

Create and optimize loops

Audio that loops seamlessly creates a perfect soundtrack for many Flash animations and video spots.

More Help topics

Create a loop

1 In the Tasks panel, click Create Loop.

   In the transport controls, Soundbooth automatically enables the Loop Playback option ⤴.

2 Press the spacebar to start playback.

3 In the Editor panel, select a time range for the loop.

4 In the Tasks panel, fine-tune the Loop In Point, Out Point, and Duration settings.

   💡 *Click the Previous Beat and Next Beat arrows to quickly adjust the selected range. Click the zoom icons to magnify related audio.*

5 Select from the following options:

   Lock Duration Restricts the loop to the currently specified duration. If you need a loop of precise length, select this option. Then manually adjust the in point and out point, or click the Previous Beat and Next Beat arrows to automatically jump from one potential loop to the next.

   Show Beat Indicators Displays beats in the Editor panel. (If beats are densely spaced, an orange gradient appears at the top of the waveform. To see individual beats, zoom in.)

6 Click Save Loop As to save to a new file.

Smooth out loop transitions

By default, Soundbooth transitions from loop out points to in points without cross-fading. That default works well for percussive sounds, but more sustained sounds benefit from a smooth cross-fade.

1 In the Create Loop task, expand the Settings section.

2 Select Auto-Smooth Loop Point.

Preview only loop transitions

By default, Soundbooth previews the entire loop. To precisely fine-tune the transition from out point to in point, preview the transition instead.

1 In the Create Loop task, expand the Settings section.

2 Select Play Transition Only.

Adjust beat detection settings

1 To display beats in the Editor panel, choose View > Beats.

2 To fine-tune beat detection, choose Tasks > Create Loop. Then, in the Tasks panel, expand the Settings section and adjust the following options:

   Minimum and Maximum Frequency Specify the main frequency range that beats occur in. The fundamental frequencies for bass drum, for example, range from 40 to 80 Hz. By contrast, snare drum emphasizes frequencies from 1,000 to 2,000 Hz.

   Sensitivity Detects main beats at low settings, and adds intermediate beats at high settings.

💡 *Appropriate frequency and sensitivity settings depend upon audio content. Experiment with different settings while viewing displayed beats in the Editor panel.*

Tempo-Based Detection Adjusts beat placement to better reflect rhythmic, musical patterns.
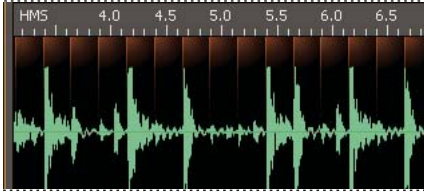
Beats in the Editor panel

💡 *To display beats as simple lines without orange semicircles, deselect Use Gradients in the Appearance preferences. (See "Change interface brightness and colors" on page 12.)*

Stretch time and shift pitch

With time-stretching, you can perfectly fit voice-overs or other audio to the length of a video scene or web animation. Independently, you can apply pitch-shifting adjustments that range from subtle enhancements to extremely high or low effects.

1 In the Editor panel, select the desired audio.

2 Choose Processes > Change Pitch And Timing.

3 Drag the Time Stretch or Pitch Shift sliders to adjust processed audio relative to existing audio. (For example, to shrink audio to half its current duration, specify a Timestretch value of 50%.)

💡 *The Current and New Duration values show length before and after time-stretching. If you need a specific length, change the New value.*

4 To more accurately adjust a solo performance, select Solo Instrument Or Voice.

5 To adjust the timbre of instruments and voices, maintaining realism during pitch shifts, select Preserve Speech Characteristics. If realism isn't desirable (for example, if you want to exaggerate the sound of pitch shifts), deselect this option.

6 Click Preview, and adjust settings as needed. To compare processed and original audio, click the power button ⏻.

7 When you achieve the best-sounding results, click OK to process the audio file.

More Help topics

"Select time ranges" on page 33

Undo and redo

Undo or redo changes

Each time you start Soundbooth, it keeps track of the edits you perform. They aren't permanently applied to a file until you save and close it, giving you unlimited undo and redo capability.

• To undo a change, choose Edit > Undo *[name of change]*. The Undo command conveniently indicates which change you're undoing (for example, Undo Delete or Undo Normalize).

• To redo a change, choose Edit > Redo *[name of change]*.

To redisplay warnings you disabled by selecting Don't Show Again, click Reset All Warning Dialogs in the General section of the Preferences dialog box.

More Help topics

Revert to history states and snapshots

While the Undo and Redo commands restrict you to an incremental sequence of changes, the History panel lets you instantly revert back to any previous change. Use the panel to quickly compare processed and original audio or discard a series of changes that produced undesired results.

History states disappear when you reopen a file, but if you save to Adobe Sound (ASND) format, you can always revert to history snapshots.

Note: When you work with very large audio files, delete unnecessary history states and snapshots to clear disk space and improve performance.

Revert to states

• To revert to any history state, click it.

• To incrementally move through states, either press the up and down arrows on the keyboard, or choose Step Backward or Step Forward from the panel menu.

Delete states

• To delete all states, choose Clear History from the panel menu.

• To delete a specific state, select it, and then click the trash icon 🗑 .

Note: Deleting history states also removes related Undo commands.

Create a snapshot

When audio content is in a state you might want to return to, create a snapshot, and save in ASND format.

1 In the History panel, click the Snapshots icon 📷 , and select New Snapshot.

2 Enter a name, and click OK.

To change a snapshot name, repeat step 1, and select Rename Snapshot.

Revert to a snapshot

❖ In the History panel, click the Snapshots icon 📷 , and select the snapshot name.

Delete a snapshot

1 In the History panel, click the Snapshots icon , and select Delete Snapshot.

2 Select the snapshot from the Name menu, and click OK.

More Help topics

"Choosing an audio file format" on page 70

"Saving and mixing down multitrack files" on page 63

Chapter 6: Effects

Effects optimize and enhance audio, giving your projects a polished, professional sound.

Applying effects

Preview and apply effects

Soundbooth lets you apply up to five effects at once, using an *effects rack*. Prior to applying a rack, you can preview it, and then customize, bypass, or remove individual effects. Bypassing effects temporarily disables them so you can compare processed and unprocessed audio. Removing effects frees up space for different effects or a new arrangement of existing ones. (Rearranging effects in a different sequence produces different sonic results.)

1 If you want to process a specific range of audio, select it in the Editor panel. (If you don't make a selection, Soundbooth processes the entire file.)

2 From the Effects menu, choose an effect. Repeat this step to add up to five effects.

💡 *The last effect you added appears at the top of the Effects menu for easy access.*

3 Click the Play button ▶ to preview processed audio.

4 In the Effects panel, do any of the following:

• To adjust an effect's settings, click Settings.

• To return to an effect's default settings, click Reset.

• To bypass a specific effect, click its Power button ⏻ .

• To bypass all effects, click the rack Power button in the lower-left corner of the panel.

• To remove specific effects, select them in the Effects panel, and click the Delete button 🗑 .

• To remove all effects, right-click the panel, and choose Remove All Effects.

5 To process the audio, click Apply To File or Apply to Selection.

💡 *To automatically display settings when you add an effect, select Auto-Open Effect Custom Settings in the General section of the Preferences dialog box.*

More Help topics

"Selecting audio" on page 33

"Effects reference" on page 47

Use effect presets

Many effects provide default *presets* with settings optimized for common tasks. In addition to these defaults, you can create custom presets for individual effects or entire racks. You access both types of presets in the Effects panel.
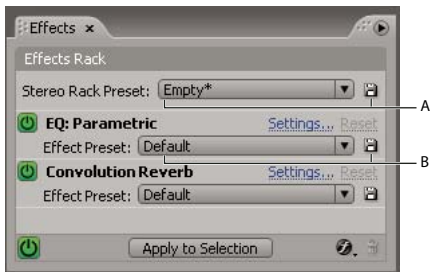
Accessing presets in the Effects panel
A. Preset options for entire rack B. Preset options for individual effect

• To load a preset, choose it from the Effect Preset or Rack Preset menu.

💡 *To re-create the settings you most recently applied, choose Last Applied Settings or Last Applied Rack.*

• To save current settings as a preset, click the Save Rack Preset or Save Effect Preset button 🖫 .

• To delete the current preset, choose Delete Preset from the Effect Preset or Rack Preset menu.

💡 *To modify an existing preset, choose it from the menu, adjust settings as desired, and then save it with the same name.*

Effects reference

About standard and advanced effects

For most effects, Soundbooth offers both standard and advanced versions. Standard effects provide basic settings, which you can quickly optimize. Advanced effects provide detailed settings, which you can fine-tune to achieve precise sonic results.

💡 *To change the sonic character of a standard effect, choose an option from the Effect Preset menu.*

More Help topics
"Applying effects" on page 46

Analog Delay effect

The Analog Delay effect creates both echoes and subtle effects. Delays of 35 milliseconds or more create discrete echoes, while delays of 15–35 milliseconds create a simple chorus or flanging effect. (However, the results won't be as effective as the Chorus/Flanger effect, because the delay settings don't change over time.) Further reducing a delay to 10–15 milliseconds adds stereo depth to a mono sound.

💡 *Make sure the audio file is long enough for the delay to end. If echoes are cut off before they fully decay, undo the effect, add several seconds of silence (Edit > Insert > Silence), and then reapply the effect.*

Standard settings
Amount Controls the ratio of original to delayed sound.

Advanced settings

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.

More Help topics

"Applying effects" on page 46

Chorus/Flanger effect

The Chorus/Flanger effect combines two popular delay-based effects. The Chorus option 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. Use this effect to enhance vocal tracks or add stereo spaciousness to mono audio.

The Flanger option creates a psychedelic, phase-shifted sound by mixing a varying, short delay with the original signal. This effect was originally created by sending an identical audio signal to two reel-to-reel tape recorders, and periodically pressing the flange of one reel to slow it down.

Standard settings

Amount Controls the intensity of the effect.

Advanced settings

Chorus Simulates several voices or instruments playing at once.

Flanger Simulates the delayed, phase-shifted sound originally heard in psychedelic music.

Speed Controls the rate at which the delay time cycles from zero to the maximum setting.

Width Specifies the maximum amount of delay.

Intensity Controls the ratio of original to processed audio.

Transience Emphasizes transients, giving them a sharper, more distinct sound.

More Help topics

"Applying effects" on page 46

Compressor effect

The Compressor effect reduces dynamic range, producing consistent volume levels and increasing perceived loudness. Compression is particularly effective for voice-overs, because it helps the speaker stand out over musical soundtracks and background audio.

For examples of highly-compressed audio, listen to recordings of modern pop music. By contrast, most jazz recordings are lightly compressed, while typical classical recordings feature no compression at all.

Standard settings
Amount Controls the level of compression.

Advanced settings
Threshold Sets the input level at which compression begins. The best setting depends on audio content and 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.

Ratio Sets a compression ratio between 1-to-1 and 30-to-1. For example, a setting of **3** outputs 1 dB for every 3-dB increase above the threshold. Typical settings range from 2 to 5; higher settings produce the extremely compressed sound often heard in pop music.

Attack Determines how quickly compression starts after audio exceeds the Threshold setting. The default, 10 milliseconds, works well for a wide range of source material. Use faster settings only for audio with quick transients, such as percussion recordings.

Release Determines how quickly compression stops when audio drops below the Threshold setting. 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 amplitude after compression. Possible values range from -30 dB to +30 dB, where 0 is unity gain.

More Help topics
"Applying effects" on page 46

Convolution Reverb effect

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.

In Soundbooth, you can use the Convolution Reverb effect to reproduce 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.

Note: Because Convolution Reverb requires significant processing, you may hear clicks or pops when previewing it on slower systems. These artifacts disappear when you apply the effect.

Standard settings
Amount Controls the ratio of original to reverberant sound.

Advanced settings
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, lusher 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.

More Help topics
"Applying effects" on page 46

Distortion effect

Use the Distortion effect to simulate blown car speakers, muffled microphones, or overdriven amplifiers.

Standard settings
Amount Controls the level of distortion.

Advanced settings
Symmetric Creates identical curves in the Positive and Negative graphs.

Positive and Negative graphs Specify separate distortion curves for positive and negative sample values. The horizontal ruler (x-axis) indicates input level in decibels; the vertical ruler (y-axis) indicates output level. The default diagonal line depicts an undistorted signal, with a one-to-one relationship between input and output values.

Click and drag to create and adjust points on the graphs. Drag points off a graph to remove them.

To copy one graph to another, click the arrow buttons between them.

Reset Returns a graph to its default, undistorted state.

Curve Smoothing Creates curved transitions between control points, sometimes producing a more natural distortion than the default linear transitions.

Time Smoothing Determines how quickly distortion reacts to changes in input levels. Level measurements are based on low-frequency content, creating softer, more musical distortion.

dB Range Changes the amplitude range of the graphs, limiting distortion to that range.

Linear Scale Changes the amplitude scales of the graphs from logarithmic decibels to normalized values.

Post-filter DC Offset Compensates for any sample offset introduced by distortion processing. Such offsets can cause audible pops and clicks when edited. (For a definition, see "DC offset" on page 78.)

More Help topics
"Applying effects" on page 46

Dynamics effect

The Dynamics effect can be used as a compressor, limiter, and 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.)

Standard settings

Amount Controls the overall level of expansion, compression, and limiting.

Advanced settings

Threshold Sets the input level at which dynamics processing begins. Note that limiting and compression begin above their threshold settings, while expansion begins below.

The best Threshold setting varies for each phase of dynamics processing:

• **Limiter** Set the threshold near the maximum amplitude you want to achieve.

• **Compressor** Set the threshold around 5 dB below the peak input level to retain more dynamic range. Or try settings around 15 dB below the peak input level to greatly reduce dynamic range.

• **Gate/Expander** Set the threshold at the level of background audio you want to remove.

Ratio Sets a compression or expansion ratio of between 1-to-1 and 30-to-1. For example, a compression setting of **3** outputs 1 dB for every 3-dB increase above the threshold. By contrast, an expansion setting of **3** reduces volume by 3 dB for every 1-dB drop below the threshold.

Typical compression settings are moderate (around 2–5), producing a natural sound. Typical limiting and expansion settings are more extreme (around 10–30). High limiting ratios ensure that audio produces maximum volume without distortion. High expansion ratios ensure that undesirable background sounds become inaudible.

Attack Determines how quickly dynamics processing starts after audio passes the threshold. The defaults (5 milliseconds for limiting, 10 milliseconds for compression and expansion) work well for a wide range of source material. Use faster settings only for audio with quick transients, such as percussion recordings.

Release Determines how quickly dynamics processing stops when audio passes the threshold. 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 amplitude after dynamics processing. Possible values range from -30 dB to +30 dB, where 0 is unity gain.

More Help topics

"Applying effects" on page 46

EQ: Graphic effect

The EQ: Graphic 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.

💡 *Use the graphic equalizer to quickly enhance voice-overs and music tracks, reproduce the muffled sound of telephone audio, and more.*

Standard settings

Low Controls a shelving EQ centered at 125 Hz.

Low-Mid Controls a peaking EQ at 500 Hz.

High-Mid Controls a peaking EQ at 2,000 Hz.

High Controls a shelving EQ at 8,000 Hz.

Advanced settings

Bands buttons Increase or decrease control by narrowing or expanding each frequency band. Options include intervals of one octave (10 Bands), one-half octave (20 Bands), or one-third octave (30 Bands).

Frequency sliders Set the amplitude level of each frequency band.

More Help topics

EQ: Parametric effect

The EQ: Parametric 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 starting around 80 Hz, and insert a 60-Hz notch filter.

💡 *Use the standard version of this effect to quickly apply preset equalization curves. Use the advanced version to precisely adjust such curves.*

Standard settings

Amount Controls the ratio of original to equalized sound.

Advanced settings

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 settings below.*

Highpass and Lowpass Enable Activate shelving filters at either end of the frequency spectrum.

Band 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.

More Help topics

"Applying effects" on page 46

Mastering effect

Mastering describes the complete process of optimizing audio files for a particular medium, such as radio, video, CD, or the web. In Soundbooth, you can quickly master audio with the 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.

Standard settings

Amount Controls the level of processing.

Advanced settings

Equalizer Adjusts the overall tonal balance. For details, see the Advanced settings for "EQ: Parametric effect" on page 52.

Reverb Adds ambience. Drag the Mix slider to change the ratio of original to reverberant sound.

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.

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.

Loudness Maximizer Increases perceived volume by limiting dynamic range and boosting output levels so that peaks approach the digital maximum of zero dBFS. As you drag the slider to the right, brickwall limiting begins at increasingly lower input levels (a setting of 100 begins limiting at –20 dBFS).

Output Gain Boosts or cuts amplitude after processing above. Possible values range from -30 dB to +15 dB, where 0 is unity gain.

More Help topics

"Repairing audio" on page 39

"Maximize volume" on page 38

Phaser effect

Similar to flanging, phasing shifts the phase of an audio signal and recombines it with the original, creating psychedelic effects first popularized by musicians of the 1960s. But unlike the Flanger effect, which uses variable delays, the Phaser effect sweeps a series of phase-shifting filters to and from an upper frequency. Phasing can dramatically alter the stereo image, creating unearthly sounds.

Standard settings

Amount Controls the ratio of original to processed sound.

Advanced settings

Stages Specifies the number of phase-shifting filters. A higher setting produces denser phasing effects.

Intensity Determines the amount of phase-shifting applied to the signal.

Depth Determines how far the filters travel below the upper frequency. Larger settings produce a wider tremolo effect.

Mod Rate Modulation rate controls how fast the filters travel to and from the upper frequency. Specify a value in Hz (cycles per second).

Phase Diff Determines the phase difference between stereo channels. Positive values start phase shifts in the left channel, negative values in the right. The maximum values of +180 and -180 degrees produce a complete difference and are sonically identical.

Upper Freq Sets the upper-most frequency from which the filters sweep. To produce the most dramatic results, select a frequency near the middle of the selected audio's range.

Feedback Feeds a percentage of the phaser output back to the input, intensifying the effect. Negative values invert phase before feeding audio back.

Mix Controls the ratio of original to processed audio.

Output Gain Adjusts the output level after processing.

More Help topics

"Waveform measurements" on page 2

"Applying effects" on page 46

Vocal Enhancer effect

The Vocal Enhancer effect quickly improves the quality of voice-over recordings. The Male and Female modes automatically reduce sibilance and plosives, as well as microphone handling noise such as low rumbles. Those modes also apply microphone modeling and compression to give vocals a characteristic radio sound. The Music mode optimizes soundtracks so they better complement a voice-over.

Male Optimizes audio for a man's voice.

Female Optimizes audio for a woman's voice.

Music Applies compression and equalization to music or background audio.

More Help topics

"Applying effects" on page 46

Chapter 7: Multitrack mixing and editing

Multitrack files let you mix together multiple audio files and scores, quickly creating sophisticated soundtracks for video and Adobe Flash.

For a video about multitrack mixing, see www.adobe.com/go/lrvid4078_sb

Creating multitrack files, tracks, and clips

Create a multitrack file

When you first open Soundbooth, it automatically creates a three-track multitrack document. You can create new multitrack documents at any time.

❖ Choose File > New, and select any of the following:

Multitrack File Creates an empty multitrack document.

Multitrack File From Video Inserts the open video file as a clip on the Video track.

Multitrack File From Waveform Inserts the open audio file as a clip on track 1.

Multitrack File From Channels Inserts each stereo channel from the open audio file as separate, mono clips on tracks 1 and 2.

To edit stereo channels separately, select Multitrack File From Channels.
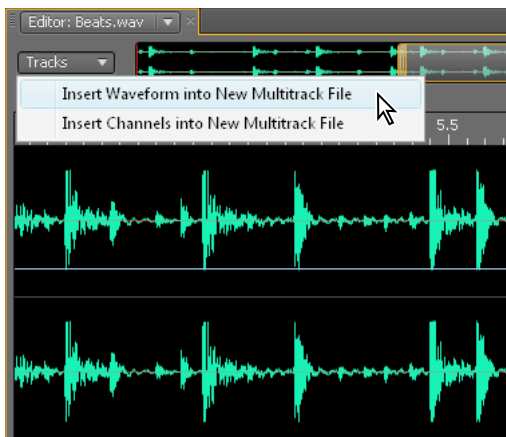
In Editor panel, Tracks menu provides alternative method for inserting open audio files into new multitrack files

More Help topics

"Open existing files" on page 16

"Record a new file" on page 17

"Saving and mixing down multitrack files" on page 63

Create, name, resize, or delete tracks

Create an audio track

1   Select the track you want the new track to appear below.

2   From the Tracks menu in the upper left of the Editor panel, select Add Audio Track.

Create a video track

You can create one video track, which appears at the top of the Editor panel.

❖   From the Tracks menu in the upper left of the Editor panel, select Add Video Track.

Name a track

❖   In the upper-left corner of the track controls, click the current name, and type a new one.
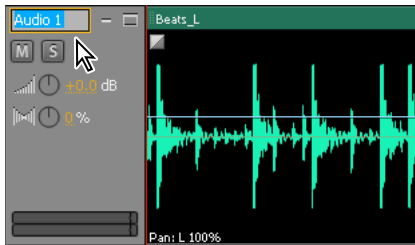
Naming a track

Resize tracks

Drag the divider between tracks up or down to change the default height of all tracks.

Delete a track

1   Select the track.

2   From the Tracks menu in the upper left of the Editor panel, select Delete Selected Track.

More Help topics

"Working with Adobe Flash and video" on page 65

Insert, copy, split, or delete clips

In a multitrack file, each audio clip corresponds to a unique audio file. If you insert the same audio file multiple times, Soundbooth copies the file for each clip. By contrast, if you insert the same video or score file multiple times, Soundbooth references one source file.

Insert a clip

❖   Do any of the following:

*   To insert a clip at a specific time point, position the current-time indicator at that point. Then choose Edit > Insert > File.

*   To insert a currently open file, drag it from the Files panel to the Editor panel.

    If you drag to an empty area below existing tracks, Soundbooth creates a track.

- To insert a sound effect from the Resource Central panel, download it, and then drag it to the Editor panel.

Copy a clip

❖ Ctrl-drag (Windows) or Command-drag (Mac OS) a clip to a new location on the multitrack.

Split a clip

❖ Select the clip, and then do any of the following:

- Choose Clip > Split Clip.

- Right-click (Windows) or Ctrl-click (Mac OS) and choose Split Clip from the context menu.

- Press Ctrl+K (Windows) or Command+K (Mac OS).

Delete a clip

❖ Select the clip, and press Delete.

More Help topics

"Position the current-time indicator" on page 25

"Open existing files" on page 16

Relinking missing video and score files

Multitrack ASND files contain all necessary audio files but link to video and score files. If a video or score file moves from its original location, the Video panel and score clip indicate that the media is offline. To relink a multitrack file to offline media, either move the media back to its original location, or complete the following procedure:

1 In the Editor panel, right-click the video or score clip, and choose Link Media.

2 Navigate to the media file, and click Open.
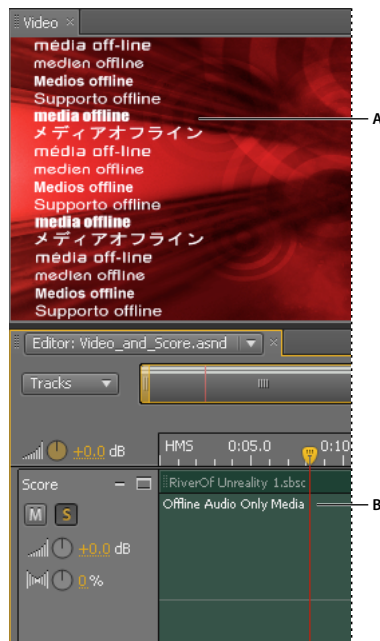
Offline media messages
A. Video panel B. Score clip

More Help topics

Mixing and editing tracks and clips

Mix tracks and clips

Master Volume and track controls
A. Master Volume B. Track Mute and Solo C. Track Volume D. Track Pan

Adjust the volume of the overall mix

❖ In the upper left of the Editor panel, drag the Master Volume knob.

Mute or solo a track

To silence an entire track, mute it. To hear a specific track, solo it.

❖ In the track controls, click the Mute or Solo button.

Adjust track volume or pan

❖ In the track controls, drag the Track Volume or Track Pan knobs.

Adjust clip volume or pan

1 In the Editor panel, select the clip, and hover the pointer over it.

2 In the controls that appear on the clip, drag the Volume ◌ or Pan ◌ knobs.

💡 *To quickly note pan percentage for any clip, look in the lower left corner.*
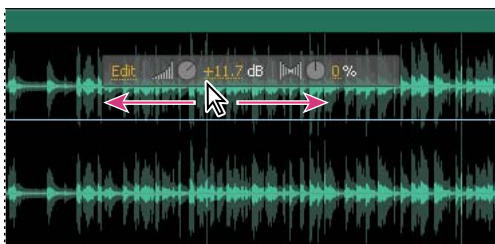
Changing volume of multitrack clip

Fade a clip

You can always readjust multitrack fades.

❖ At the far left or right of the clip, drag the Fade In ◩ or Fade Out ◪ handle.

For information about similar controls for individual files, see "Precisely fade in or out" on page 36.

More Help topics

"Revert to history states and snapshots" on page 44

"Saving and mixing down multitrack files" on page 63

Automate mixes with keyframes

To create dynamic, automated mixes, use *keyframes*. Keyframes change audio volume and score parameters over time. To gradually increase volume, for example, add a –3 dB keyframe at the beginning of a clip and a +3 dB keyframe at the end. Soundbooth automatically calculates, or *interpolates*, all the intermediate values using one of two transition methods:

- *Hold* ☐ transitions create an abrupt change in value at each new keyframe.

- *Linear* ◇ transitions create a smooth, gradual change between keyframes.
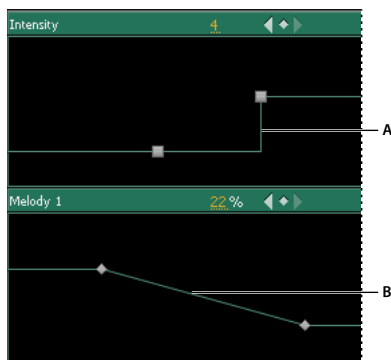
Transitions between keyframes
A. Hold B. Linear

Add a keyframe

Do either of the following:

- Position the pointer over a parameter line. When a plus sign ▸₊ appears, click.

- Position the current-time indicator where you'd like a score parameter to change. Then click the Add Keyframe icon ◈ for the parameter, or change the parameter value.

Navigate between score keyframes

❖ In the Editor panel, click the Previous Keyframe ◀ or Next Keyframe ▶ icon for a score parameter.

Display different score parameters

❖ Select options from the Parameters menu at the top of a score clip.

💡 To quickly adjust score parameters, see "Shortcuts for score keyframes" on page 76.

Select all keyframes for a property

❖ Right-click any keyframe for the property, and choose Select All Keyframes.

Adjust keyframe values

❖ Drag a keyframe up or down.

Adjust keyframe time position

❖ Drag a keyframe right or left.

Change the transition between two keyframes

❖ Right-click the first keyframe, and choose either Linear to gradually transition from one value to the next, or Hold to abruptly change values.

Delete keyframes

❖ Right-click a keyframe, and choose Delete Selected Keyframes.

More Help topics

Move and trim clips in the multitrack display

Move a clip

❖ Drag it left or right. Or, to place a clip more precisely, enter a specific Start or End Time in the Properties panel (Window > Properties).

By default, clips snap to edges of clips on other tracks. To temporarily disable snapping while moving or trimming clips, hold down Ctrl (Windows) or Command (Mac OS). To permanently disable snapping, deselect View > Snapping > Enabled.

Trim a clip

You can readjust the trim points for multitrack clips at any time.

❖ Position the pointer over the left or right edge of the clip, and drag.

To trim a clip with maximum precision, double-click it to enter single-clip editing mode, and drag the trim handles at either end of the file.

More Help topics

Edit and process a single multitrack clip

To add effects or process specific time or frequency ranges, use the single-clip editing mode. This mode expands the view of a clip and enables the Solo button [S] for the containing track. (To hear other tracks for context, deselect the Solo button.)

1 Double-click the clip, or click Edit in the upper-right corner.

2 If you want to process a specific range, select it.

3 Modify the audio with commands in the Edit, Processes, or Effects menus.

> 💡 *Editing or processing audio is permanent, unless you create a snapshot. (See "Revert to history states and snapshots" on page 44.) You can readjust effects, however, unless you apply them to audio.*

4 To return to the multitrack display, double-click the clip header, or click Back in the upper-right corner.

More Help topics

"Editing and repairing audio files" on page 28

"Edit audio from video in a multitrack file" on page 68

"Preview and apply effects" on page 46

Customizing scores

With scores, anyone can create compelling music and atmospheric soundtracks for video and Adobe Flash. Choose from dozens of score templates in a wide variety of styles, and then customize compositions for specific projects. Automate parameters such as Intensity and Melody with keyframes, creating dynamic soundtracks that evolve over time.

Install and preview scores

Install scores

Soundbooth installs two scores by default. To access additional scores, do the following:

❖ In the Scores panel, click More Scores At Resource Central.

> 💡 *If you need to change the location for the two default scores, click Browse in the General section of the Preferences dialog box. Then move the Adobe Soundbooth Scores folder to the new location.*

Preview scores

Score template (SBST) files determine score parts and parameters. You can preview these files in the Scores panel or Adobe Bridge.

1 In the Scores panel, select an option from Genre menu to filter the list.

2 Select a score, and then click the Start Score Preview icon ▷.

> 💡 *To automatically play selected scores, click the Auto Play Toggle icon ◀�))*.

More Help topics

"Browsing assets with Adobe Bridge" on page 16

"Open existing files" on page 16

Insert a score in a multitrack file

Quickly insert a score

❖ Drag it from the Scores or Resource Central panel to the Editor panel.

Note: Before you can drag scores from the Resource Central panel, you must download them.

Precisely insert a score

❖ In the Editor panel, position the current-time indicator. In the Scores panel, select a score, and click the Add Score ⌁ icon.

Replace an existing score

❖ In the Editor panel, select the existing score clip. In the Scores panel, select a new score, and click the Replace Score icon ⟳.
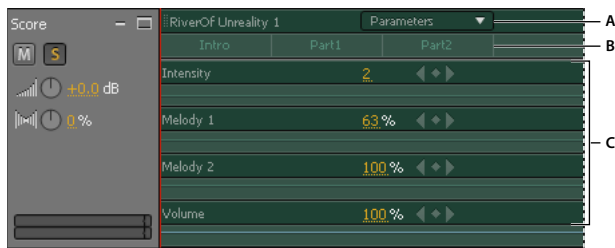
Score clip in Editor panel
A. Menu determines displayed parameters B. Score parts C. Adjustable parameter lines

💡 *By default, the Properties panel appears when you select a score clip in the Editor panel. If you prefer to retain the currently displayed panels, deselect Auto-Open Properties Panel in the General section of the Preferences dialog box.*

More Help topics

"Position the current-time indicator" on page 25

Adjust score duration and parts

1 In the Editor panel, select a score clip. Then do either of the following in the Properties panel:

• Choose a standard duration such as 30 seconds from the Variation menu. Standard durations arrange score parts in an ideal order set by the template creator.

• Enter a custom Clip Duration. To optimize the arrangement of musical parts for the new duration, choose Automatic from the Variation menu.

2 Choose an Intro/Outro option to add or remove these parts. (If editing these parts changes score length too drastically, readjust the Clip Duration setting.)

💡 *In the Editor panel, you can visually adjust start time, end time, and duration by dragging the middle or edge of a score clip.*
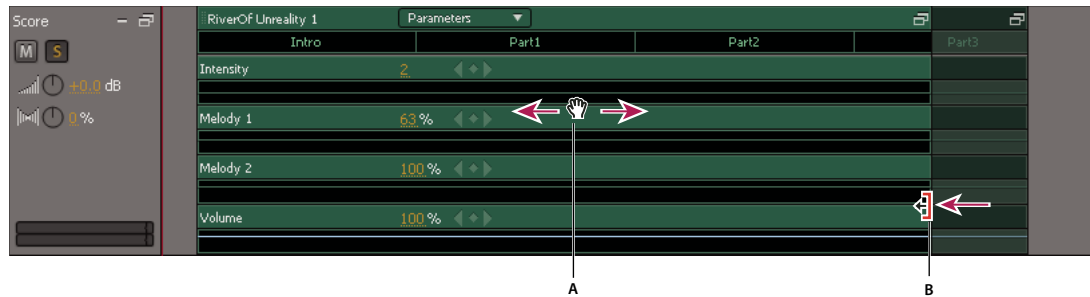
Changing score start time (A) and duration (B) in the Editor panel

More Help topics

"Playing audio" on page 24

Set static score parameters

By default, you automate score parameters. (See "Automate mixes with keyframes" on page 59.) But if you prefer static parameter values for a particular score, disable keyframe automation.

1 In the Editor panel, select the score clip.

2 In the Properties panel, click Basic.

3 Do either of the following:

  • To quickly change the overall character of the score, choose an option from the Preset menu.

  • To fine-tune the score, drag sliders for individual parameters such as Intensity or Melody. (Intensity adjusts the mix from sparse to dense instrumentation. Additional parameters adjust the level of specific elements.)

  💡 *You can also drag parameter lines in the Editor panel.*

Saving and mixing down multitrack files

Saving multitrack files

Soundbooth saves multitrack files in Adobe Sound (ASND) format, which stores history snapshots and lets you readjust trimming and effects. (See "Revert to history states and snapshots" on page 44.) If you're working with Adobe Premiere Pro, After Effects, and Flash, import ASND files in those applications for maximum versatility. When you edit and resave a multitrack mix in Soundbooth, the related ASND file automatically updates in other applications.

Mix down complete multitrack files

❖ To mix down to common formats like mp3 or WAV, which a broad range of applications support, choose File > Export > Multitrack Mixdown.

Mix down individual clips

If you want to import an individual multitrack clip in another application, mix the clip down. The exported clip reflects all edits and processing applied in Soundbooth.

1 In the Editor panel, select the clip.

2 Choose File > Export > Clip Mixdown

More Help topics

Chapter 8: Working with Adobe Flash and video

Soundbooth tightly integrates with Adobe Flash and video. Intuitive visual tools let you quickly repair a variety of common audio problems, from wireless crackle to unwanted background noise. And Soundbooth Scores give you the power to create a customized musical or atmospheric soundtrack.

Working with Flash cue points

Together with ActionScript®, Flash cue points trigger events or navigate to different points in time. Both Soundbooth and the Adobe Media Encoder let you preview video and insert cue points. However, only Soundbooth pairs that functionality with powerful audio editing features, including a waveform display that helps you place cue points with maximum precision.

Cue point workflow

In Soundbooth, Flash cue points are called *markers*. To precisely place markers in media files and convert them into cue points that interact with ActionScript, follow this workflow:

1. Create ActionScript for the cue points

For more information, see Learning ActionScript 3.0.

2. Edit an audio, video, or multitrack file in Soundbooth

Sources can include files exported from other applications, or files you create in Soundbooth. (See "Opening, creating, and recording files" on page 16.)

3. Add markers, and set their cue point properties

In the file, insert and adjust markers. (See "Use markers" on page 32.) Then, in the Markers panel, set cue point type (event or navigation), and enter name and value parameters.

4. Output cue points for Flash projects

To embed cue points into media files, save to FLV format. (See "Save entire files or selected ranges" on page 69.) Or, to edit cue points separately from media files, export to XML format, and then import the XML in either Adobe Flash Professional or the Adobe Media Encoder.

Set properties for cue points

1 In the Markers panel, select a marker, and expand the Marker Details section.

2 For Type, select either Event or Navigation.

3 Adjust parameters:

- To add new parameters, click the Add Parameter button . Enter a name and value, and click OK.

- To edit existing parameters, double-click them.

- To delete parameters, select them, and click the Clear Parameter button ⊑ .

Export or import cue points in XML files

To edit cue points in Adobe Flash Professional or the Adobe Media Encoder, export them to XML. Import the resulting XML file in Soundbooth if you want to adjust the corresponding markers in media files.

Export markers as cue points

❖ Choose File > Export > Markers.

To embed Flash cue points in media files, save them to FLV format. (See "Save entire files or selected ranges" on page 69.)

Export spoken words as cue points

To trigger animated events based on speech, export each spoken word as a cue point.

1 Create a text transcript. (See "Analyze speech for text XMP metadata" on page 22.)

2 Choose File > Export > Speech Transcription.

Import cue point XML files

❖ Choose File > Import > Markers.

More Help topics

"Use markers" on page 32

Edit audio from Adobe Flash, Premiere Pro, or After Effects

From Adobe Flash Professional, Adobe Premiere Pro, or After Effects, you can quickly access Soundbooth to edit, repair, and enhance audio files. While these applications support many standard audio formats, the most versatile choice is Adobe Sound (ASND). That format lets you readjust fades and effects, save snapshots that restore previous edits, and create multitrack mixes.

1 In Adobe Flash, Premiere Pro, or After Effects, select an audio or ASND file.

2 Choose Edit > Edit In Adobe Soundbooth

The file opens in Soundbooth. When you save your edits in that application, the file updates automatically in Adobe Flash and video applications.

For more information, search for "Soundbooth" in Adobe Flash, Premiere Pro, or After Effects Help.

More Help topics

"Choosing an audio file format" on page 70

"Revert to history states and snapshots" on page 44

"Multitrack mixing and editing" on page 55

Working with video files

Maintaining video synchronization when editing audio

When you edit audio from a video file, trim handles and the Cut, Delete, and Crop commands silence audio without changing file length. This approach maintains synchronization, so you can easily incorporate edited audio into your video projects.

Trim handles silence audio, maintaining video synchronization. (The Cut, Delete, and Crop commands do the same.)

By default, the Video panel appears when you open a video file. If you find video distracting, deselect Auto-Open Video Panel in the General section of the Preferences dialog box. Alternatively, to display the Video panel more prominently, choose Window > Workspace > Edit Audio to Video.

More Help topics

"Choosing a video file format" on page 71

"Trim, crop, or delete audio" on page 35

Edit audio from video in a multitrack file

When you insert a video clip in a multitrack file, audio adjustments are limited to track volume and pan settings. To access the full range of editing options, extract the video soundtrack to a separate audio clip.

❖ In the Editor panel, click Extract Audio in the upper right of the video clip.

Soundbooth creates an audio clip on a new track and enables the Mute button on the video track.
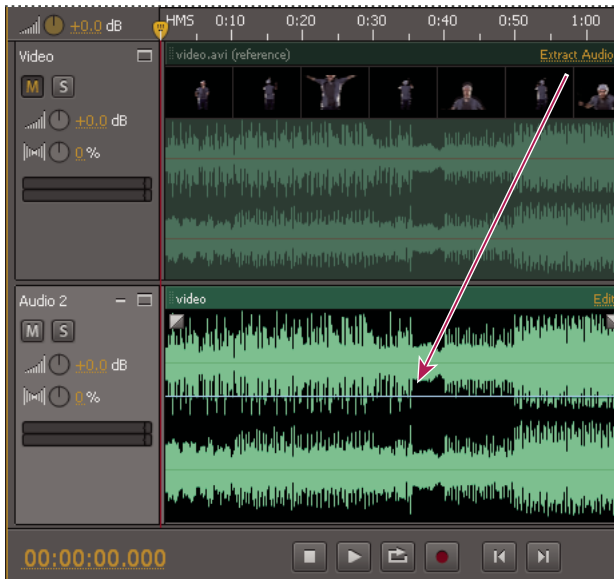
To edit audio from a video clip, extract it.

More Help topics

"Multitrack mixing and editing" on page 55

Chapter 9: Saving audio and video files

Adobe Soundbooth offers a comprehensive set of file formats that support a wide variety of output types, ranging from the web to high-definition video.

Saving and closing files

Save entire files or selected ranges

Use the File > Save command to either save new files or quickly resave existing files and retain their settings. Use File > Save As commands to save selections or change settings such as file format, sample rate, and bit depth.

💡 *To exclude metadata from files created with Save As commands, deselect Include Source XMP Metadata in the Metadata section of the Preferences dialog box.*

1 Choose one of the following:

- File > Save. (If you're resaving, skip remaining steps.)

- File > Save As.

- File > Save Selection As.

2 Specify a file location, name, and format. Then click Save. (If you're saving to ASND format, skip the next step.)

3 Set format-specific options. For more information, see one of the following:

- "Options for standard audio formats" on page 72.

- "Options for video and AAC formats" on page 73.

💡 *To save all open files, choose File > Save All. For existing files, Soundbooth retains current settings. For any new files, however, save options appear.*

More Help topics

"Choosing file formats" on page 70

"Saving and mixing down multitrack files" on page 63

"Export or import cue points in XML files" on page 66

Export audio channels to mono files

If you want to edit or output individual channels of a stereo or surround-sound file, export them to mono files. Soundbooth appends the channel name to each exported file name and saves in WAV format. For example, a stereo source file named Jazz.aif produces mono files named Jazz_L.wav and Jazz_R.wav.

1 Choose File > Export > Channels To Mono Files.

2 Specify a location for the files, and click OK.

Soundbooth automatically opens each exported file.

💡 *To convert stereo channels into separate multitrack clips, choose File > New > Multitrack File From Channels.*

More Help topics

"Saving and mixing down multitrack files" on page 63

"Export or import cue points in XML files" on page 66

Close files

- To close the current file, choose File > Close.

- To close all open files, choose File > Close All.

- To close files in the Files panel, select them, and click the Close File button 🗑 .

To quickly close selected files, press Delete or Backspace.

Choosing file formats

Choosing an audio file format

For maximum flexibility, perform all editing in Adobe Sound (.asnd) format. If you primarily use Soundbooth with Adobe Flash, Premiere Pro, and After Effects, use ASND format exclusively. When you want to share files with a broader range of applications, such as disc-burning utilities, save to the WAV or AIF formats. Save to the compressed mp3 and AAC formats only when creating files for the web or portable media players.

Here are details about each audio format:

Adobe Sound (.asnd)

ASND format lets you readjust fades and effects, and store history snapshots that restore previous edits. This format also supports multitrack mixes, which let you combine multiple audio files and scores. Adobe Flash, Premiere Pro, and After Effects import ASND files, tightly integrating those applications with Soundbooth.

Advanced Audio Coding (.aac)

AAC is a compressed format based on MPEG-4. The alternative filename extension .m4a appears in many applications, including Apple iTunes. At similar bit rates, this format achieves higher audio quality than mp3 format.

Audio Interchange File Format (.aif)

AIF is the standard, uncompressed audio file format for Mac OS.

mp3 Audio (.mp3)

mp3 is the most widely used format for web-based audio and portable media players. This format highly compresses file size, optimizing audio for fast downloads. The compression process, however, slightly reduces quality and introduces artifacts, particularly in quiet passages.

If you edit an mp3 file, avoid saving it to the same format. Resaving in mp3 format recompresses data, lowering audio quality.

Windows Waveform (.wav)

Windows Waveform is the standard, uncompressed audio format for the Windows operating system.

More Help topics

"Revert to history states and snapshots" on page 44

"Saving and mixing down multitrack files" on page 63

"Options for standard audio formats" on page 72

Choosing a video file format

For many video projects, you'll save ASND or WAV audio files from Soundbooth, and then combine those with related video in an application like Adobe Premiere Pro. If the flexibility of separate audio and video files isn't necessary, however, save to a video format directly from Soundbooth. Choose a format optimized for the output type:

•   For hard disk playback, choose uncompressed AVI or QuickTime.

•   For the web, choose compressed Flash Video, MPEG1, QuickTime, or Windows Media.

•   For standard resolution DVDs, choose MPEG-2-DVD.

•   For high-definition video, choose H.264 (MPEG4).

Here are details about each video format:

Adobe Flash Video (.flv)
FLV format lets you present video in Adobe Flash Player, a free, widely available browser plug-in. Adobe Flash Player can play either standalone FLV files, or those you embed into Flash animations in SWF format.

Microsoft AVI(.avi)
The Windows version of Soundbooth supports AVI, Microsoft's standard video file format. AVI is a container format that supports both uncompressed video and a variety of codecs.

MPEG-1 (.mpg)
The Windows version of Soundbooth supports MPEG-1, a compressed format commonly used on CD-ROM and the web. This format produces picture quality comparable to VHS at quarter-screen frame size.

Note: Due to their keyframe-based compression and variable bitrate, MPEG formats require significant processing power and time to generate.

MPEG-2 (.mpg, .m2v)
MPEG-2 format delivers SVHS picture quality, much higher than MPEG-1. A variation of this format is part of the original DVD specification, but MPEG-2 is also supported by the Blu-ray and HD-DVD standards.

MPEG-4 (.mp4, .m4v)
Maintains the same perceived quality level as MPEG-2, but offers better compression, reduced file size, and support for interactivity. The H.264 codec is supported by the Blu-ray Disc and HD-DVD formats.

QuickTime (.mov)
QuickTime is the standard video format for Mac OS, though its use isn't limited to that platform. However, Windows users must install QuickTime for Windows to view and work with MOV files. In addition to full-resolution video, QuickTime supports streaming video and many different types of compression.

Windows Media Video (.wmv)

The Windows version of Soundbooth supports Windows Media Video, Microsoft's format for compressed, streaming video. Typically, WMV files are viewed in Windows Media Player, but other applications and plug-ins also support this format.

Though the formats above are primarily used for video output, you can specify audio-only output in the Export Settings dialog box.

More Help topics

"Options for video and AAC formats" on page 73

Options for standard audio formats

ASND options

In the Adobe Sound Documents section of the Preferences dialog box, the following settings apply to newly created files. (To change these settings for an existing ASND file, open it, and choose File > Adobe Sound Document Settings.)

Multitrack Time Display For new multitrack files, determines the time range initially visible in the Editor panel.

Embedded Mixdown Determines the audio quality of ASND files in other Adobe applications. Either select Flash Compatible, or select Custom and specify a sample rate and bit depth appropriate for your video projects. (See "Understanding sample rate" on page 4 and "Understanding bit depth" on page 5.)

More Help topics

"Saving and mixing down multitrack files" on page 63

WAV and AIF options

When you save in WAV or AIF format, the Save As Options dialog box appears. This dialog box provides a compact group of settings that you can quickly optimize.

General options
File Type Indicates the file format.

Range Specifies whether to save the entire file or the currently selected range.

Add To Files Panel Adds the saved file to the Files panel.

Audio options
Compressor Specifies the audio codec, if any. The codecs available depend on the platform and file format. In Mac OS, only the Uncompressed option is available. In Windows, WAV format provides these additional options:

• **IMA ADPCM** Compresses 16-bit audio to 4-bit. This compression scheme can be a good alternative to MPEG-1; it provides fairly fast decoding and degrades audio quality only slightly.

• **Microsoft ADPCM** Compresses 16-bit audio to 4-bit using a slower method than IMA ADPCM, but quality can be superior, depending upon audio content.

• **CCITT A-Law** Compresses 16-bit audio to 8-bit, but produces dynamic range equivalent to13 bits (78 dB). This compression scheme is the standard for European telephone systems.

- **CCITT u-Law** An older version of CCITT compression that requires slightly more processing than A-Law. This compression scheme is the standard for North American and Japanese telephone systems.

- **GSM 6.1** Compresses audio at a ratio of 10:1. Outside of Japan, GSM is the most popular worldwide standard for mobile phones.

- **DSP GroupTrueSpeech (Windows XP only)** A competing standard for mobile phones, this scheme compresses audio at a ratio of 15:1.

Sample Rate Sets the rate at which audio is digitally sampled. Higher sampling rates support a wider frequency range, increasing audio quality and file size. To determine the proper setting for common output types, see "Understanding sample rate" on page 4.

Note: Though you can work with any sample rate in Soundbooth, your sound card may not be capable of playing all rates properly. To determine supported sample rates, consult the documentation for the card.

Sample Type Sets the audio bit depth. Higher bit depths provide more dynamic range and reduce distortion, though they increase file size and processing time. To determine the proper setting for common output types, see "Understanding bit depth" on page 5.

Channels Specifies whether the saved file contains one mono channel, two stereo channels, or six (5.1) surround channels. You can save to an equal or lesser number of channels. (For example, you can save a stereo file to mono, but not a mono file to stereo.)

mp3 options

When you save in .mp3 format, Soundbooth provides the following options:

Bitrate Determines the rate in kilobits per second. Higher rates increase file size and audio quality; lower rates reduce size and quality.

To hear how the selected bitrate affects audio quality, click the Preview button ▶.

Channels Encodes audio either in mono, which reduces file size, or stereo, which increases file size but retains spatial information.

Options for video and AAC formats

When you save in AAC, Adobe Flash Video, AVI, MPEG, QuickTime, or Windows Media format, the Export Settings dialog box appears. This dialog box lets you access the wide array of export options provided by the Adobe Media Encoder. For detailed information, click the Help link below.

Chapter 10: Keyboard shortcuts

Keyboard shortcuts help you speed up the editing process.

Finding and customizing shortcuts

Find shortcuts

❖ To find shortcuts, do any of the following:

• For menu commands, look to the right of command names.

• For tools, look to the right of tool tips. (To display tool tips, hold the pointer over a tool.)

• For a complete list of shortcuts, choose Edit > Keyboard Shortcuts.

Customize shortcuts

You can customize nearly all default keyboard shortcuts and add shortcuts for other commands.

1 Choose Edit > Keyboard Shortcuts.

2 In the Command column, select the command you want to customize.

💡 *To quickly switch between custom panel layouts, assign shortcuts to the Workspace commands.*

3 If you want to replace or remove an existing shortcut, choose it from the Shortcuts For Command menu.

4 Do any of the following:

• To create a shortcut, click inside the Press Shortcut box, and press the desired key combination. Then click Assign.

• To remove a shortcut, click Remove.

Note: If you enter a key combination that's already in use, Soundbooth displays an alert. Click Yes to transfer the shortcut to a different command, or No to retain the existing assignment.

5 When you finish making changes, click OK.

Save or delete custom sets of shortcuts

1 Choose Edit > Keyboard Shortcuts.

2 Do either of the following:

• To save a custom set, click Save As, enter a name, and click OK.

• To delete a custom set, choose it from the Set menu, and then click Delete.

Restore default shortcuts

1 Choose Edit > Keyboard Shortcuts.

2 From the Set menu, choose Default Set.

Common shortcuts

Shortcuts for playing and zooming audio

This partial list includes the playing and zooming shortcuts that many Soundbooth experts find most useful.

Result	Windows shortcut	Mac OS shortcut
Start and stop playback	Spacebar	Spacebar
Enable or disable looped playback	Ctrl+L	Command+L
Shuttle playback left or right (press again to increase speed)	J or L	J or L
Stop shuttled playback	K	K
Select Hand tool	H	H
Select Zoom tool	Z	Z
Move current-time indicator to beginning of file	Home	Home
Move current-time indicator to end of file	End	End
Zoom in	=	=
Zoom out	–	–
Zoom out fully	\	\
Zoom to selection	Shift+S	Shift+S
Zoom to selection in point	Shift+Q	Shift+Q
Zoom to selection out point	Shift+W	Shift+W

You can also zoom in and out with the mouse wheel.

More Help topics

"Playing audio" on page 24

"Zoom audio" on page 30

Shortcuts for editing audio

This partial list includes the editing shortcuts that many Soundbooth experts find most useful.

Result	Windows shortcut	Mac OS shortcut
Expand Editor panel to fill application window.	Tilde (~) key. Do not press Shift.	Tilde (~) key. Do not press Shift.
Select Time Selection tool	T	T
Select Frequency Selection tool	F	F
Select Rectangular Marquee tool	R	R
Select Lasso tool	G	G
Auto-heal an audio glitch	Ctrl+U	Command+U

Result	Windows shortcut	Mac OS shortcut
Insert silence	Ctrl+T	Command+T
Insert marker	M or * (asterisk key)	M or *(asterisk key)
Add most recent effect to rack	Ctrl+Shift+E	Command+Shift+E

More Help topics

"Selecting audio" on page 33

"Repairing audio" on page 39

"Applying effects" on page 46

Shortcuts for score keyframes

This partial list includes the keyframe shortcuts that many Soundbooth experts find most useful.

To access these shortcuts, click a score parameter lane.

Result	Windows shortcut	Mac OS shortcut
Select previous or next keyframe	Left or right arrow	Left or right arrow
Select multiple keyframes	Shift+left or right arrow	Shift+left or right arrow
Adjust parameter values	Up or down arrow	Up or down arrow
Adjust time position	Alt+left or right arrow	Option+left or right arrow

More Help topics

"Automate mixes with keyframes" on page 59

"Customizing scores" on page 61

Chapter 11: Digital audio glossary

The glossary is your guide to unfamiliar terms in common audio workflows and multiple Adobe Soundbooth features. If you don't find a term here, search for it in Help to find a feature-specific definition.

Common audio terms

A

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 80.) 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 78.)

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.

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 The number of bits used to represent audio amplitude. Higher bit depths provide greater dynamic range but increase file size. For details, see "Understanding bit depth" on page 5.

C

clipping In digital audio, distortion that occurs when the amplitude of a signal exceeds the maximum level for the current bit depth. 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" on page 79.)

crossfade A fade from one audio track to another.

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.

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.

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" on page 77 and "decibel (dB)" on page 78.)

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 the Fourier theory, which allows quick analysis of frequency and amplitude. In Soundbooth, you adjust spectral resolution by specifying different FFT sample sizes. (See "Customize the spectral display" on page 30.)

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 2.)

H

hertz (Hz) Cycles per second. A unit of measurement that describes the frequency of a sound. (See "frequency" on page 79.)

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 or Core Audio 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.

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. (To master audio in Soundbooth, see "Mastering effect" on page 53.)

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 78.)

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. The Normalize feature in Soundbooth normalizes audio to 0.3 dBFS, avoiding clipping while ensuring optimal volume.

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.

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 Soundbooth 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 2.)

processes Intensive functions that require dedicated computer power, briefly preventing you from editing audio. Processes permanently change audio data. (Compare with "real time" on page 80.)

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 Describes functions that immediately respond to user input. Real time functions leave audio data intact, so you can readjust settings in the future. (Compare with "processes" on page 80.) In Soundbooth, real time functions include effects and mixing controls.

resample To convert a sound file to a different sample rate and bit depth. In Soundbooth, you resample files by using the Save As command. (See "Save entire files or selected ranges" on page 69.)

reverb The reverberant sound produced by an acoustic space, such as a room or concert hall. Reverb consists of dense, discrete echoes that arrive at the ear so rapidly that the ear can't separate them.

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

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 provide a greater frequency range but increase file size. For details, see "Understanding sample rate" on page 4.

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 signal-to-noise ratio is around 92 dB.

sound card A hardware device that lets your computer play and record audio. Sound cards can be built into the computer motherboard, added as an internal card, or connected via an external USB or FireWire device.

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 Soundbooth, 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.

T

tempo The rhythmic speed of music, normally measured in beats per minute. (See "beats per minute (bpm)" on page 77.)

W

waveform The visual representation of an audio signal, displayed as amplitude across time in Soundbooth. (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*.