

Live 7 Audio Fact Sheet

Much of Ableton's recent development effort has been focused on carefully and objectively testing Live's fundamental audio performance. As a result of this testing, we have implemented a number of low-level improvements to the audio engine. We have also written this fact sheet to help users understand exactly how their audio is (or is not) being modified when using certain features in Live that are often misunderstood, as well as tips for achieving the highest quality results.

As mentioned above, the focus of our research has been on objective (that is, quantifiable and measurable) behavior. We make no claims about what you can *hear* because we can't possibly predict the variables that make up your listening environment, audio hardware, hearing sensitivity, etc. Furthermore, this paper makes no claims about how Live compares to other audio software. Rather, it is a summary of measurable facts about what Live 7 actually *does* under various conditions.

Testing and Methodology

As of this writing, every version of Live is subjected to a suite of 473 automated tests that cover every aspect of Live's functionality. We add additional tests as we add features, and we will never release an update unless it passes every test.

Neutral Operations

Procedures in Live that will cause absolutely no change in audio quality are referred to as *neutral operations*. You can be sure that using these functions will never cause any signal degradation. Applying neutral operations to audio that was recorded into Live ensures that the audio will be unchanged from the point of analog-to-digital conversion. Applying neutral

operations to files imported into Live ensures that the imported audio will be identical to the files saved on disk. Applying neutral operations to files being exported from Live ensures that the quality of your output file will be at least as high as what you heard during playback.

The list of neutral operations found below is provided primarily as an abstract reference; while all of these operations are, in fact, neutral, it is important to remember that each of them may (and almost certainly will) occur within a context that also contains non-neutral operations. For example, running an audio signal through an effects device is a non-neutral operation. So any neutral operations that occur after it will, of course, still result in audio that is altered in some way. Even a gain change is, technically, non-neutral.

Neutral operations include:

Undithered Rendering

The Export Audio/Video command renders Live's audio output to a file on disk. Rendering is a neutral operation under certain conditions:

- the sample rate of the rendered file is the same as that set for the audio hardware in Live's Preferences.
- no non-neutral operations have been applied.

Live's rendering performance is tested by loading three types of unprocessed audio files (white noise, fixed-frequency sine waves and sine sweeps) in 16-, 24- and 32-bit word lengths and rendering these to output files, also with varying bit resolutions. Phase cancellation testing of the original and output files shows the following:

- rendering to a file with the same bit depth as the original results in complete phase cancellation.
- rendering to a file with a higher bit depth than the original results in complete phase cancellation.
- rendering to a file with a lower bit depth than the original results in the smallest amount of distortion possible within a 32-bit system.

Matching sample rate/no transposition

Playback of an unstretched audio file in Live is a neutral operation, provided that the file's sample rate is the same as that set in Live's Preferences and that the file is played back without transposition. This is verified by cancellation tests of rendered output. Please note that "playback" in this context refers only to the audio within Live, prior to the point at which it reaches your audio hardware.

Unstretched Beats/Tones/Texture/Re-Pitch Warping

If the tempo of a Clip is the same as the tempo of the Set, that clip will play back unstretched. In this case, if the Warp mode of the Clip is set to Beats, Tones, Texture or Re-Pitch (but not Complex), playback will be neutral. Any Warping caused by changing the Set's tempo is non-permanent, and audio that plays back unwarped at a given tempo will always play back unwarped at that tempo, even if the tempo is changed and then changed back. For example, if you've recorded some tracks at 120 BPM, but then decide you'd like to slow the tempo down to record a particularly difficult solo passage, the original tracks will play back neutrally again after returning the tempo to 120 BPM. Only the recording made at the slower tempo will be stretched. Please note that the Global Groove functionality works by modifying the positions of Warp markers. This means that playback of audio clips with Global Groove applied will be non-neutral even at the original tempo.

The neutrality of unstretched clip playback is verified by performing cancellation tests on rendered output.

Summing at Single Mix Points

Live 7 uses double precision (64-bit) summing at all points where signals are mixed, including Clip and return track inputs, the Master track and Racks. Mixing in Live is thus a neutral operation for signals mixed at any single summing point. This is tested by loading pairs of 24-bit files (white noise and fixed-frequency sine waves and their phase-inverted complements), adding the pairs together eight times and rendering the output as 32-bit files. All tests result in perfect phase cancellation.

Please note that, while 64-bit summing is applied to each *single* mix point, Live's internal processing is still done at 32-bit. Thus, signals that are mixed across multiple summing points may still result in an extremely small amount of signal degradation. This combination of 64-bit summing within a 32-bit architecture strikes an ideal balance between audio quality and CPU/memory consumption.

Recording external signals (bit depth \geq A/D converter)

Recording audio signals into Live is a neutral operation, provided that the bit depth set in Live's Preferences window is the same or higher than that of the A/D converters used for the recording. In this context, "neutral" means "identical to the audio as it was delivered to Live by the A/D converters."

Recording internal sources at 32 bit

Audio that is recorded via internal routing will be identical to the source audio, provided that the recording was made at 32 bits. To ensure neutral recordings of plugin instruments and any audio signals that are being processed by effects plugins, internal recording at 32 bits is recommended. Please note, however, that if the source audio is already at a lower bit depth, internal recording at that bit depth will also be neutral (assuming that no effects are used); internally recording an unprocessed 16 bit audio file at 32 bits will not increase the sound quality.

The neutrality of internal recording is verified using cancellation tests.

Freeze, Flatten

When tracks are frozen, the audio files that are created are 32 bit, which ensures that they will not be lower quality than the audio heard prior to freezing. But there are some special cases involving Freeze that result in non-neutral behavior and should be noted:

- Frozen Arrangement View tracks can include audio material that extends beyond the end of the clip itself, such as reverb tails and delay repetitions. Frozen Session View

tracks, however, are always exactly two loop cycles long, so any audio that extends beyond two loop cycles during un-frozen playback will be cut off after freezing.

- Time-based effects like reverbs and delays are processed in realtime for unfrozen clips, so stopping playback during a reverb or delay tail will allow the tail to continue. In contrast, frozen tails are rendered as audio, and so will stop abruptly during playback.
- Any parameter automations are rendered as part of the audio file for frozen Arrangement View clips. Frozen Session View clips, however, take a “snapshot” of all parameter values at the Arranger’s 1.1.1 position and retain them for the duration of the frozen clip. This is analogous to the behavior with unfrozen clips; when playing normal clips in Session View, any Arrangement automations are “punched out” until the Back to Arrangement button is pressed.
- Frozen clips are always played back with Warp on and in Beats mode, which means they are subject to the same non-neutral behavior as any other Warped audio files.
- Any devices with random parameters (e.g., the Chance control in the Beat Repeat device) will no longer exhibit random behavior after freezing. This is because, as with time-based effects, the random values that were in place at the moment of freezing will be rendered as part of the new file, and will thus no longer be calculated in real-time.

Please note that the Flatten command replaces any original clips and devices with the audio files created by freezing. When using this command, it is important to keep in mind the special cases above – what you hear after freezing is exactly what you will get when flattening, so if the results are not to your liking, be sure to unfreeze and make any necessary changes to device parameters before invoking the Flatten command.

This procedure is tested by rendering the output of an audio track and comparing it to the frozen audio from the same track via phase cancellation to ensure that the files are identical.

Bypassed Effects

Bypassed effects in Live are removed from the signal flow. This is true for both Live’s built-in effects devices and third-party VST and AU plugins. Consequently, audio at the output of a bypassed effect is identical to the audio at the input. Please note, however, that effects devices with parameters that inherently require delay (e.g., the Look Ahead settings in

Compressor) will still introduce this delay when bypassed, in order to maintain automatic delay compensation with the rest of the project. In most cases, the effects of this behavior will be completely inaudible.

The neutrality of bypassed effects is tested by loading one instance of each of Live's effects devices into an audio track, deactivating them, and then rendering the output of the track. The rendered file is then compared to the rendered output of the same track with no loaded devices. Phase cancellation testing of the two files confirms that they are identical.

Routing

The routing of signals within Live is a neutral operation. The signal at the routing destination will be identical to the signal at the routing source. It is important to note that Live's flexible routing architecture allows for a variety of scenarios, including routing from before or after any track's effects or mixer and tapping the output of individual sample slots within the Impulse instrument. In these cases, it is likely that the signal heard at the output point will be different from the signal heard prior to routing, because it has been tapped before reaching the end of its original signal chain.

Splitting Clips

Clips which are already neutral will remain neutral after splitting. Splitting only affects playback position within the sample, and has no effect on the sample data itself. Playback across a split boundary is seamless and sample-accurate.

The neutrality of clip splitting is tested under a variety of conditions:

- splitting unwarped clips with loop on and off;
- splitting warped but unstretched clips with loop on and off;

In all cases, output is rendered and compared with the output of an unsplit version of the same source. Phase cancellation testing of the two files confirms that they are identical.

Non-Neutral Operations

Procedures in Live that will cause a change in audio quality are referred to as *non-neutral operations*. Users can be guaranteed that using these operations will cause at least some change to the signal. Applying non-neutral operations to files imported into Live ensures that the imported audio will differ from the files saved on disk. Applying non-neutral operations to files being exported from Live ensures that what you hear during realtime playback will be different from what will end up in your new file.

Non-neutral operations include:

Playback in Complex Mode

The algorithm used in Complex Warp Mode is an entirely different technology from the algorithms behind Beats, Tones and Texture modes. Although Complex Mode may sound better, particularly when used with mixed sound files that are comprised of many different kinds of audio material, it is never neutral – not even at the original tempo. Because of this, and because of the algorithm’s increased CPU demands, we recommend using it only in cases where the other Warp modes don’t produce sufficient results.

Sample rate conversion/transposition

Sample rate conversion (during both real-time playback and rendering) is a non-neutral operation. Playback of audio files at a sample rate that is different from the rate set in Live’s Preferences window will cause signal degradation. Transposition is also a form of sample-rate conversion, and thus also results in non-neutral behavior.

To minimize potential negative results, it is recommended to do sample rate conversion as an offline process in another application. Once the samples have been converted to the sample rate that you plan to use in Live, the files can be imported without any loss of quality.

Rendering audio from Live with a sampling rate other than the one that was used while working on the project is also a non-neutral operation, and may result in a loss of sound quality. It is recommended to always render using the original sampling rate, and then

convert the rendered file using a dedicated mastering application that is optimized for these kinds of CPU-intensive, offline tasks.

While we recommend that you use a high-quality offline tool for sample rate conversion, we recognize that one of Live's core features is its ability to pitch-shift and warp audio in real time. For this situation, it is necessary to make a trade-off between CPU performance and precision. We recommend the use of the Hi-Q button for any clips which undergo transposition in a given Set. The algorithm behind the Hi-Q switch has been rewritten for Live 7, and now results in considerably lower distortion than in previous versions.

Volume Automation

Automation of volume level results in a change in gain, which is necessarily a non-neutral operation. But certain implementations of automation envelopes can result in audible artifacts, particularly if the envelopes are not calculated at a fast enough rate. In Live 7, volume automation curves are updated for each audio sample, resulting in extremely low levels of distortion.

Dithering

Whenever rendering audio to a lower bit depth, it is a good idea to apply dithering in order to minimize artifacts. Dithering (a kind of very low-level noise) is inherently a non-neutral procedure, but it is a necessary evil when lowering the bit resolution.

Please note that Live's internal signal processing is all 32-bit, so applying even a single gain change makes the resulting audio 32-bit as well – even if the original audio is 16- or 24-bit. Dither should never be applied more than once to any given audio file, so unless you are mastering and finalizing in Live, it is best to always render at 32-bit and avoid dithering altogether.

Recording external signals (bit depth < A/D converter)

Recording audio signals into Live is a non-neutral operation if the bit depth set in Live's

Preferences window is lower than that of the A/D converters used for the recording. This is not recommended.

Recording internal sources below 32 bit

Audio that is recorded via internal routing will lose quality if the recording is made at a bit depth below 32 bits. To ensure neutral recordings of plugin instruments and any audio signals that are being processed by effects plugins, internal recording at 32 bits is recommended. Please note, however, that if the source audio is already at a lower bit depth, internal recording at that bit depth will also be neutral (assuming that no effects are used); internally recording an unprocessed 16 bit audio file at 32 bits will not increase the sound quality.

Consolidate

Consolidating clips in the Arrangement View creates new audio files, which are non-neutral in comparison to the original audio data. Specifically, the new files will be normalized, with their clip volumes adjusted to play back at the same volume as heard prior to consolidation. Normalization is a gain change, which is a non-neutral operation. Also, the new files will be created at the sample rate and bit depth set in Live's Preferences window, which may differ from those in the original audio files.

Clip fades

With Clip Fade enabled, a short (up to 4 ms) fade is applied to the clip start and end to avoid clicks at the clip edges. This is a non-neutral operation.

Panning

Live uses constant power panning with sinusoidal gain curves. Output is 0 dB at the center position and signals panned fully left or right will be increased by +3 dB. In order to

minimize this volume change, it may be helpful to narrow the overall stereo width before doing extreme panning. This can be done via the Width control in the Utility device.

Global Groove Changes

Under most conditions, playback of a warped clip that is at the same tempo as the Set is a neutral operation. However, if the Global Groove amount is adjusted and the clip has a Clip Groove setting other than Straight, playback will be non-neutral at any tempo.

Tips for Achieving Optimal Sound Quality in Live

For users looking to achieve optimal audio quality in Live, we have provided a list of recommended practices and program settings.

- Decide which sample rate to use for a project prior to beginning work, rather than changing the sample rate while working on the project.
- Record audio into Live using high-quality hardware components (audio interface, cables, etc.) and at the highest sample rate and bit depth your interface and computer will support.
- Avoid using samples that are at different sample rates within the same project. If you want to work with such files, we recommend that you first convert them to the sample rate set for your audio interface in an offline application that is optimized for this task.
- For all audio clips, disable both the Warp and Fade options in the Clip View.
- Do not adjust the Transpose and Detune controls for any clips.
- Always render at 32-bit and at the sample rate set for your audio interface. If you need audio files at a different sample rate and/or bit depth, we recommend that you convert your rendered files in an offline application that is optimized for these tasks, rather than in Live.

Please note that these practices, while ensuring optimal audio quality, disable some of Live's functionality – in particular, stretching and synchronization.

Summary and Conclusions

Ableton wrote this paper in order to help users understand exactly how audio is affected when performing various procedures in Live. Our focus has been on functions that have proven over the years to cause confusion or uncertainty, and the list of both neutral and non-neutral operations presented here is necessarily incomplete.

We encourage you to refer to this paper if you have any questions about the functions discussed, but we also encourage you to [contact us](#)¹ if you have additional questions that we haven't answered. Live is a product, but it is also an ongoing process, and one that benefits greatly from user input.

¹contact@ableton.com