



Apple Digital Masters: Music as the Artist and Sound Engineer Intended

Whether you are a major label, an indie, an aggregator, a recording artist, or an audio engineer, you provide the most important ingredient for Apple Music and the iTunes Store — the music itself. It is our job to faithfully and accurately deliver this music to fans around the world exactly as you intend it to be heard.

We've designed workflows and protocols to facilitate the best possible results, ones that live up to your highest standards for audio quality. To achieve this transparency, you can use tools and technologies from Apple to ensure delivery of the highest-quality master recordings possible into our ecosystem.

With more than a billion active iOS devices, Mac computers, Apple TVs and HomePods in the world, there's never been a better time to refine, codify, and communicate updated information to you about the best workflows and processes to produce the billions of files streamed and delivered daily to our mutual customers in over 100 countries around the world.

Note: **Apple Digital Masters** replaces **Mastered for iTunes** as the name of the program to reflect the fact that these audio advantages are available to our entire music catalog across the Apple ecosystem, whether streamed or downloaded. All Mastered for iTunes releases are now badged as Apple Digital Masters and the Mastered for iTunes software tools are still usable for Apple Digital Masters creation.

Innovation and Excellence in Sound

Apple celebrates a rich history and tradition of innovation and excellence in sound for computing as well as content creation. The original Mac was engineered fully capable of supporting audio without additional hardware or software, making it one of the first personal computers ever to ship with sound. In fact, Apple received a GRAMMY Award® for technical excellence in music, the first and only such award ever given to a personal computing company.

When iTunes launched, the decision was made to standardize on AAC instead of the more popular MP3 format simply because AAC clearly provides superior audio quality compared to other codecs at similar bit rates. We worked with Dolby and Fraunhofer and have constantly iterated the Apple AAC encoder to reach the industry-leading quality it achieves today. And today Apple Music delivers lossless versions of these masters to our subscribers using the Apple Lossless Audio CODEC (ALAC.)

Following the guidelines and protocols outlined in this document to create and audition AAC encodes, you can achieve simply amazing results: dynamic range that's superior to red book CD audio, and a final product that's virtually indistinguishable from the original studio recording.

Mastering for Digital Delivery

For decades, the standard for consumer digital audio was the compact disc (CD) and most mastering was done with CDs in mind. But with billions of streams and downloads encoded as AAC to date worldwide, AAC is now the standard for digital music. It only makes sense to create masters specifically for this format.

What is AAC?

AAC (Advanced Audio Coding) is a format for compressing and encoding digital audio. AAC achieves the portability and convenience of compressed and encoded digital audio while retaining audio quality that's indistinguishable from much larger digital files.

The iTunes catalog was initially offered in 2003 as 128 kbps AAC files, many of which were encoded from original CD masters. They sounded great—in fact, these downloads led the industry in sound quality. More than 100 million songs were sold in this format in a little over a year, changing the landscape of legal digital music forever.

But innovation didn't stop there. We have iterated on our encoder for over fifteen years. Our AAC encoder is now able to transparently encode higher resolution audio—creating files that retain the small footprint, portability, and ease of use Apple is known for. (And they sound amazing.)

High-Resolution Digital Recording

Digital audio, such as that on a CD, generally uses Linear Pulse Code Modulation (LPCM — often referred to simply as PCM) to represent audio signals. LPCM works by taking snapshots of the analog audio signal and assigning each a numerical value.

The resolution of an LPCM recording is determined by the sample rate (how many times per second samples are taken) and the bit depth (how many bits are used to represent each sample). Higher sample rates can capture higher frequencies, and higher bit depths can accurately represent a greater dynamic range.

The standard for CDs is 16-bit 44.1kHz resolution, meaning that the analog signal is sampled 44,100 times per second and each sample is given a value between -32,768 and 32,767. This resolution is often referred to as 44/16.

The Nyquist sampling theorem states that to accurately represent a signal, one must use a sampling rate double that of the highest frequency being represented. The highest frequency audible to humans is around 20kHz; therefore a sampling rate of over 40kHz is required to accurately capture the audible range of frequencies. Compact discs' 44.1kHz rate is adequate for this need.

Even so, many experts feel that using higher sample rates during production provides better-quality audio and a superior listening experience in the end product. For this reason, higher sample rates of 48, 88.2, 96 and even 192 kHz have become standard. In addition, 24-bit recording has become the norm—increasing workable headroom and total dynamic range.

Challenges with Encoding Higher-Resolution Recording

An inherent challenge of working with high-resolution audio has been that both the sample rate and the bit depth must be reduced to match the specifications used in mainstream distribution, whether for CD or AAC. This can be done either with software (“in the box”) or by running the mix through an additional analog stage and resampling. There are pros and cons with each technique, but both can add noise and/or distortion.

Downsampling, as the word implies, is the process of using Sample Rate Conversion

Frequency is the number of vibrations per second and is measured in Hertz (Hz). Human hearing spans a range from roughly 20Hz to 20kHz.

Bit rate is different than **bit depth**. Bit rate indicates how much data is being used per second and is calculated using the sample rate and bit depth. A 256 kbps AAC file is encoded with a target bit rate of 256 kilobits/second. It utilizes Variable Bit Rate (VBR) encoding, which uses each bit strategically, dynamically allocating less data for simple sections and more data for complex passages.

Dynamic Range, when used as a general audio term, refers to the range of possible volumes—the difference between the softest and the loudest parts.

In digital audio, **aliasing** refers to audible artifacts created when higher frequencies are sampled at an insufficient rate. The result is distortion. A visual metaphor for aliasing can be found in the “wagon wheel effect”—a rapidly spinning wagon wheel filmed at a low frame rate can appear to be moving backward.

Clipping is a form of audio distortion and can be caused in many ways. In general, it is the result of the amplitude of a signal becoming too great to be accurately represented by a system. In an amplifier, this can occur if one attempts too much amplification—the top of the signal is cut off, or “clipped.” In digital audio, this can occur when a signal falls outside of an allowed bit- depth range.

(SRC) to lower the sample rate (for example, from 96kHz to 48kHz). This process may create aliasing, an undesired effect.

Dithering (or adding dither) is a technique used when reducing a file’s bit depth (for example, from 24-bit to 16-bit). It is an attempt to reduce the distortion inherent in this process. Dither is a tradeoff—the distortion is reduced at the cost of added noise. (Another bit depth reduction option is called “truncating,” where the additional bits are simply removed. This can cause quantization distortion.)

Improved Conversion and Encoding for AAC

Apple’s latest encoding methodology is a two-step process. The first step in the encoding path is to use state-of-the-art mastering-quality Sample Rate Conversion (SRC) to resample the master file to a production sample rate.

This SRC outputs a 32-bit floating-point file which can preserve values that might otherwise fall outside of the permitted amplitude range. This file is saved in a CAF container. This critical intermediary step prevents any aliasing or clipping that could otherwise occur in SRC. It is this 32-bit floating file that’s used as the input to the encoder—this is one key reason we can achieve superior results.

Our encoders then use every bit of resolution available, preserving all the dynamic range of the 32-bit source file and eliminating the need for adding dither. The advantage of this is twofold. Not only does it obviate the need of adding dither noise, it also lets the encoders work more efficiently as they don’t need to waste resources encoding this unwanted and unnecessary noise.

By using this highly accurate file directly from our SRC and taking advantage of its clean signal, our encoder can deliver the final product exactly as the artist and sound engineers intended it to sound.

Lossless Distribution

In addition, Apple Music can now deliver lossless encodings to customers using the Apple Lossless Audio CODEC (ALAC.)

Best Practices for Apple Digital Masters

Our latest high-resolution encoding process ensures that your music is transparently and faithfully distributed in the way you intended it to be heard. However, before you submit songs to Apple for encoding, there are some best practices you can follow to ensure that your audio is optimized for Apple Music.

Provide High-Resolution Masters

Some mastering engineers prefer to control the SRC process by sending already converted files, however we ask that you deliver the highest native sample rate available. As technology advances and bandwidth, storage, battery life, and processor power increase, keeping the highest-resolution masters available in our systems allows for full advantage of future improvements to your or your client’s music.

Also, though it may not be apparent because there may not always be a physical, tangible master created in LP or CD format, the Apple Music and iTunes catalog forms an important part of the world’s historical and cultural record. These masters matter.

Master for Apple's AAC Encoder

Because 256 kbps AAC is a highly portable format, its files have the potential to be listened to in a wide range of different settings. So, while one listener may be using AirPods while riding in a loud subway car, another may wind up listening intently to a Bach cantata on a HomePod in their bedroom—or on a Denon receiver using AirPlay 2 in a home media room. A college student may be deep into Miles Davis' *Sketches of Spain* while sporting Beats Studio headphones in the campus library.

You're being provided with all the tools you'll need to encode your masters using the same software that Apple Music and the iTunes Store does so that you can audition exactly what they'll sound like to your listeners.

Loudness, Dynamic Range, and Clipping

Whether you're mastering a whisper-quiet zen flute tone poem or a heavy metal guitar fest, volume and loudness are key issues. The main tools used in mastering—equalization, compression, limiting, or combinations of these—are all different ways of controlling aspects of volume. Making decisions about gain levels, dynamic range, and frequency response is what mastering is all about.

Analog masters traditionally have volume levels set as high as possible, just shy of oversaturation, to improve the signal-to-noise ratio (SNR). With digital masters, the goal is to achieve the highest gain possible without losing information about the original file due to clipping.

With digital files, there's a limit to how loud you can make a track: 0dBFS. Trying to increase a track's overall loudness beyond this point results in distortion caused by clipping and a loss in dynamic range. The quietest parts of a song increase in volume, yet the louder parts don't gain loudness and can distort due to the upper limits of the digital format.

The advent of digital look-ahead “brickwall” limiters have made it possible to increase perceived volume even further.

Some artists and producers feel that louder is better. The trend for louder music has resulted in both ardent fans of high volumes and backlash from audiophiles, a controversy known as “the loudness wars.” While some feel that overly loud mastering ruins music by not giving it room to breathe, others feel that the aesthetic of loudness can be an appropriate artistic choice for particular songs or albums.

A less obvious issue when setting gain for digital masters can occur on playback. Whether it's a compressed file like an AAC file or an uncompressed file such as a CD, digital data goes through several processes to be converted to an analog signal for playback.

One common process is called oversampling. This upsamples the digital data at four times the original sample rate to improve the quality of the digital audio signal being converted to analog. If the original digital audio data is at 0dBFS, oversampling can result in undesirable clipping. And if the original was already clipped, oversampling can make it worse. Our recommendation is to leave at least 1 dB of headroom in order to avoid such clipping.

In both digital and analog, the highest possible level will vary from track to track,

Equalization (EQ) controls the volume of specific frequencies—such as bass and treble. There are many types of EQ, all of which control the volume of specific frequencies.

Compression automatically controls volume over time. A compressor decreases dynamic range by decreasing loud signals, increasing lower signals, or by doing both. (Audio compression shouldn't be confused with data compression, which is used to reduce digital file sizes.)

Limiting is a fast-acting form of audio compression with a high ratio—it's often used to attenuate peak levels of a signal that could otherwise result in distortion. By raising the total volume and limiting the peaks, the total dynamic range is compressed, increasing loudness.

Multi-band Compression is a combination of EQ and compression where a signal is divided into bands of frequencies which can then be isolated and compressed without affecting the other frequencies. For example, one can compress the bass without affecting the vocal.

0dBFS (or “Zero Decibels relative to Full Scale”) refers to the maximum possible values that can be represented by a digital signal. Values that are beyond the maximum value cause clipping. For example, in a 16-bit file the highest value that can be represented is 32,767, and the lowest value that can be represented is -32,768. If you attempt to represent a series of values above 32,767, they will all be clipped at 32,767 and will create artifacts that sound unpleasant and dissonant when played back.

depending on the material being mastered. Your decision about the volume and loudness of your tracks is a technical and creative choice. You might decide to take the listener on a dynamic journey through an album as a complete work, raising and lowering the volume level across the sequence of tracks to increase the music’s emotional impact. Alternately, you might pursue the loudest possible signal at all times.

Whatever you decide—exquisitely overdriven and loud or exquisitely nuanced and tasteful—we will be sure to encode it and reproduce it accurately. We ask that you avoid clipping the signal, and we have provided tools to make it easier than ever to provide clean masters that result in distortion-free streamed and downloaded music.

Loudness Normalization

Sound Check is a feature available in Apple Music and iTunes that lets listeners hear all their songs at approximately the same loudness. It first determines the loudness of a track and then stores that information in the file’s metadata. The metadata is then used to raise or lower the volume of each track to prevent jarring volume changes while a device is shuffling songs.

Many radio and streaming services use technology similar to Sound Check to control changes in loudness when playing songs on the air. MP3s have a similar technology to control volume changes, called Replay Gain. The International Telecommunication Union (ITU) is taking a similar approach in its standardization of the characterization of volume in its broadcasting standards (specifically BS. 1770).

(Sound Check can also correctly set the volume per album, rather than per song, allowing albums that rely on volume differences between tracks, such as ***The Dark Side of the Moon*** by Pink Floyd, to maintain their intended volumes.)

The effect of Sound Check, as well as other volume-controlling technologies, is that songs that have been mastered loud will be played back at a lower volume which can make tracks actually sound weaker.

Because many such technologies are available to listeners, you should always mix and master your tracks in a way that captures your intended sound, regardless of playback volume.

Remastering and Apple Digital Masters

When the CD was first becoming a popular format, many older recordings were rushed to market. To make these recordings available as soon as possible, corners were sometimes cut. In several instances, the wrong master tapes were used or the CD was mastered poorly. Many of these mistakes have since been identified and corrected, but even so, a number of record labels are remastering older material for rerelease in high-resolution formats like DVD-A or SACD.

Advances in AAC encoding provide an opportunity for anyone maintaining a catalog of older content to create a digital archive or improve an existing one. There have also been advances in digital mastering tools and techniques. In addition to higher resolutions, other technologies (such as noise reduction and pop and click removal) have undergone significant improvements. Now that there’s a clear distribution path that can take full advantage of high-resolution masters, an older recording may merit being given new life. Delicate and artful remastering can let works of genius be experienced as they were meant to be and haven’t been for years.

When remastering for Apple Digital Masters, procedures and best practices similar to mastering for new releases should be followed. You should always work from the best available master, carefully monitor gain to avoid clipping, and audition the encoding process.

Previewing the Converted Audio File

The AAC-encoded M4A file (final.m4a) created by afconvert can be previewed in any application that plays M4A files (such as the Music App or QuickTime Player). However, if you want to do a more detailed examination (such as loop sections, compare sections with the original, and so forth), you can decode the AAC-encoded data back to uncompressed LPCM data, just as it's decoded during playback. You can use afconvert to decode the AAC data to uncompressed LPCM and store it to a WAV file.

This WAV file should only be used for making a more detailed comparison or evaluation beyond what you can do by just previewing the file. **The deliverable for Apple Digital Masters is the original 24 bit PCM file.**

Use afconvert to decode the AAC-encoded data, as it will correctly decode the AAC file, taking care of some details that other applications don't do correctly. For example, the resulting decode.wav file will have the same number of audio samples as the original source.wav file that you started with, and thus it can be exactly aligned with your source file for critical listening, comparison, and evaluation.

Using afconvert to Decode Your AAC-Encoded M4A File

Use the following command to decode your AAC-encoded M4A file. In this command, **the M4A file with AAC-encoded data that was originally generated using afconvert is assumed to be called *final.m4a*; the WAV file generated from decoding the AAC-encoded data is called *decode.wav*.**

1. In Terminal, type on one line:


```
afconvert final.m4a decode.wav -d LEI24 -f WAVE
```

afclip

You can use another simple command-line tool, afclip, to check any file for clipping. This tool works by examining an audio file and identifying areas where clipping has occurred.

It accepts audio files as input, and outputs a stereo sound file containing the left channel of the original file and a right channel with graphically represented impulses corresponding to each clipped sample in the original. This sound file can then be loaded into a digital audio workstation (DAW), such as Logic, providing a visual map to locate any clipping that may have occurred.

The following image is an example of an audio file that has been loaded into a DAW to help visually locate clipping.



For more information about afclip, type `afclip -h` on the command line in the Terminal application on Mac OS X.

Using afclip to Check for Clipping

To check an audio file for clipping with afclip:

1. Open a Terminal window.
2. In the Terminal window, type the following on one line, followed by a space:
`afclip`
3. Drag and drop the audio file you wish to check onto the Terminal window.
Press Return to run afclip.

By default, afclip will give a readout of any clipping found. If there's clipping, it will also output a .wav file in the active directory.

Reading the afclip Readout

If any clipping is found, a readout will be printed to the terminal window with details for each instance of clipping.

SECONDS	SAMPLE	CHAN	VALUE	DECIBELS
14.595941	643681.00	0	-1.002716	0.023562
14.595964	643682.00	0	-1.016706	0.143904
14.595986	643683.00	0	-1.031721	0.271245
14.596009	643684.00	0	-1.027497	0.235609
14.642409	645730.25	1	1.000831	0.007211
14.642511	645734.75	1	1.003081	0.026721
14.642517	645735.00	1	1.008520	0.073693

The readout will contain the following information for each instance of clipping:

- Seconds.** The time, in seconds, where the clipping occurs.
- Sample.** The sample number that was clipped.
- Channel.** The channel of the clipped sample. A value of 0 means the clipping occurred on the left channel, while a value of 1 means the clipping occurred on the right channel.
- Value.** The raw value of the clipped sample. Since clipping happens when a value exceeds the range of -1 through 1, these values will be below -1 or above 1.
- Decibels.** The number of decibels by which the sample exceeds the clipping point.

The readout will end with a summary of how many total clipped samples the audio file contains for the left channel and the right channel.

186.064853	8205460.00	0	-1.062570	0.527150
186.064875	8205461.00	0	-1.003081	0.026723
186.312120	8216364.50	1	-1.009396	0.081235
186.312132	8216365.00	1	-1.034017	0.290550
186.312154	8216366.00	1	-1.037741	0.321778
188.297455	8303917.75	1	1.000005	0.000041
188.297460	8303918.00	1	1.008717	0.075388
188.297483	8303919.00	1	1.025134	0.215617
total clipped samples Left on-sample: 10661 inter-sample: 2876				
total clipped samples Right on-sample: 9690 inter-sample: 2487				

Some readouts will show sample values with decimals, such as 8216364.50 in the table above. This indicates that the waveform will clip between two samples of the original

audio data. This is called inter-sample clipping. Most DACs will upsample audio while converting the digital signal to an analog signal, which is known as the reconstruction process. This process can result in clipped values that aren't actually clipped in the original material, causing an uncompressed digital signal to cause clipping in playback even if all the values in the waveform don't show on-sample clipping. For this reason, it's important to check for inter-sample clipping. By default, afclip does inter-sample clip detection and includes the total number of inter-sample clipping at the bottom of the readout, next to the total number of on-sample clipping.

AURoundTripAAC Audio Unit

The AURoundTripAAC Audio Unit lets you compare audio encoded using 256 kbps AAC against a source audio file in realtime. It also includes clip and peak detection, as well as a simple listening test environment. The audio unit can be used in any audio unit host application, such as Logic or AU Lab.

Audio to WAVE Droplets

The Audio to WAVE Droplet automates the creation of WAVE (Waveform Audio File) files from any audio file (such as MPEG or CAF files) natively supported on macOS. For example, the Droplet can be used to decode AAC files to 24-bit WAVE files, automating the decoding process used by the command line utility, afconvert, and allowing you to decode multiple files at once.

To use the Audio to WAVE Droplet, drag and drop source audio files, or folders containing those files, onto the droplet. The droplet will then convert those files to WAVE format files. The created WAVE files will be titled using the names of their corresponding source files, and will be placed in the same folder as the source files.

NOTE: The files created with these droplets should be used for auditioning and comparison purposes only. **The deliverable for Apple Digital Masters is the original 24 bit PCM file.**

Addendum: Lossless Audio and Apple Digital Masters

Lossless Audio

Apple Music now offers subscribers lossless streaming and downloads using the Apple Lossless Audio CODEC (ALAC.) There is no need to send two masters—the same Apple Digital Master is used for the ALAC encodes as for the AAC encodes.