



Mastered for iTunes: Music as the Artist and Sound Engineer Intended

Whether you're a major label or a small indie, you provide the most important ingredient for iTunes—the music itself. It's our job to faithfully and accurately deliver your songs and albums to fans around the world exactly as you intend them to be heard. We've designed our tools to facilitate the best possible results, ones that live up to your highest standards for music available on the iTunes Store. To achieve this transparency, you need tools and technologies from us to ensure delivery of the highest quality master recordings possible into our ecosystem. With over 315 million iOS devices capable of playing your music, there's never been a better time than now for us to communicate, codify, and distribute updated information to you about the best tools and processes used to produce the millions of AAC files delivered daily to our mutual customers in over 50 countries around the world.

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 2002, even before the launch of the iTunes Store, 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. In working with Dolby and Fraunhofer, there have since been further improvements to AAC to get it to the level of excellence experienced on iTunes today. If you follow the guidelines outlined in this document and audition sample AAC encodes on Apple devices, you can achieve dynamic range that's superior to red book audio and a final product that's virtually indistinguishable from the original recording.

Mastering for Digital Delivery

Digital distribution is no longer an afterthought. It is today's dominant medium for consuming music and as such needs to be treated with utmost care and attention. For decades, the standard for consumer digital audio has been the compact disc, and most mastering has been done with CDs in mind.

In recent years the quality of digital music delivery has vastly increased, as has the number of digital music sales, with iTunes being a key driver of those sales. With more than 16 billion downloads encoded as AAC to date worldwide, AAC is the new standard for digital music. It only makes sense to create masters specifically for this format.

What is AAC and iTunes Plus?

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 larger digital files, such as audio from CDs.

The iTunes catalog was initially offered in 2003 as 128 kbps AAC files, many of which were encoded from the 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. Recently, using the most advanced AAC encoder, the iTunes catalog was upgraded to iTunes Plus: a variable bit rate (VBR) 256 kbps AAC encoding format. iTunes AAC encoders are now able to transparently encode high definition audio, creating files that retain the small footprint, portability, and ease of use iTunes 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 resolution PCM files during production provides better-quality audio and a superior listening experience in the end product. For this reason, 96/24 resolution is quickly becoming a standard format in the industry, and it's also common to see higher resolution files, such as 192/24.

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, such as 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 (SRC) to lower the sample rate (for example, from 96kHz to 44.1kHz). This process commonly creates aliasing, an undesired effect.

Dithering (or dithering down) is a technique used when reducing a file's bit depth (for example, from 24-bit to 16-bit). It's an attempt to reduce the distortion inherent in this process. Dithering 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

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. An iTunes Plus AAC file uses a sample rate of 44.1kHz and 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.

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 sample rate of 44.1kHz.

Because this SRC outputs a 32-bit floating-point file, it can preserve values that might otherwise fall outside of the permitted amplitude range. 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 and is one key reason for such stunning results.

Our encoders then use every bit of resolution available, preserving all the dynamic range of a 24-bit source file and eliminating the need for dithering. 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.

Best Practices for Mastering for iTunes

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 iTunes for encoding, there are some best practices you can follow to ensure that your audio is optimized for iTunes.

Provide High Resolution Masters

To take best advantage of our latest encoders send us the highest resolution master file possible, appropriate to the medium and the project.

An ideal master will have 24-bit 96kHz resolution. These files contain more detail from which our encoders can create more accurate encodes. However, any resolution above 16-bit 44.1kHz, including sample rates of 48kHz, 88.2kHz, 96kHz, and 192kHz, will benefit from our encoding process.

Don't upsample files to a higher resolution than their original format. Upsampling won't recover or add information to an audio file. Don't provide files that have been downsampled and dithered for a CD. This degrades the file's audio quality.

As technology advances and bandwidth, storage, battery life, and processor power increase, keeping the highest quality masters available in our systems allows for full advantage of future improvements to your 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 iTunes catalog forms an important part of the world's historical and cultural record. These masters matter—especially given the move into the cloud on post-PC devices.

Master for iTunes Plus

When creating a master, mastering engineers take into account the limitations and characteristics of the medium or destination format, as well as the listening environment of their audience. For example, a master created for vinyl is unlikely to be listened to in an airplane or car, and therefore is often mastered for a listening environment where a listener can hear and appreciate a wider dynamic range. Similarly, a master created for a club environment might take into account the noisiness of the intended listening environment.

Because iTunes Plus 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 white earbuds while riding in a loud subway car, another may wind up listening intently to a Bach cantata on AirPlay-equipped Bowers and Wilkins speakers or on a similarly equipped Denon receiver in a home media room. Just as likely, a college student may be deep into Miles Davis' *Sketches of Spain* while sporting Dre Beats headphones in the campus library. Keep in mind that Apple has sold more than 250 million iOS devices, and that many, many people around the world are listening to music on their iPods, iPhones, or iPads.

You're being provided with all the tools you'll need to encode your masters precisely the same way the iTunes Store does so that you can audition exactly what they'll sound like as iTunes Plus AAC files.

Further, to ensure that your audience is hearing your intended sound, Apple recommends listening to your masters on the devices your audience will be using. While this doesn't have to be done in real time, you should be confident that the sound will hold up when heard on its intended listening device and in its intended environment.

Be Aware of Dynamic Range and Clipping

Whether you're mastering a whisper-quiet zen flute tone poem or a heavy metal guitar fest, volume is a key issue. 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.

Many 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." This is solely an issue with music. Movies, for example, have very detailed standards for the final mastering volume of a film's soundtrack. The music world doesn't have any such standard, and in recent years the de facto process has been to make masters as loud as possible. 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.

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 due to the upper limits of the digital format.

Although iTunes doesn't reject files for a specific number of clips, tracks which have audible clipping will not be badged or marketed as Mastered for iTunes.

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

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. An extreme form of limiting is "brick wall limiting."

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.

make it worse. A growing consensus is emerging that digital masters should have a small amount of headroom (roughly 1dB) in order to avoid such clipping.

In both digital and analog, the highest possible level will vary from track to track, 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 only ask that you avoid clipping the signal.

Master for Sound Check and Other Volume Controlling Technology

Sound Check is a feature in iTunes and all iOS devices that lets listeners hear all their songs at approximately the same volume. It first determines the loudness of a track and then stores that information in the file’s metadata. (Songs downloaded from the iTunes Store already contain this information.) 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.

Radio uses technology similar to Sound Check to control changes in volume 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 to be too loud will be played back at a lower volume, letting listeners more easily notice any artifacts or unintentional distortion.

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 for iTunes

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.

The recent advances in encoding files for iTunes Plus provides an opportunity for anyone maintaining a catalog of older content to create a digital archive or improve an existing one. In the last five years, 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.

Materials already being remastered in high resolution formats, such as DVD-A or SACD, are excellent candidates for being remastered for iTunes. Projects that can benefit from remastering include CDs that were poorly mastered or mastered from the wrong tapes, projects that could benefit from cutting-edge noise reduction or similar modern technologies, and cases where the physical media is in danger of deteriorating.

Unlike physical media, digital archiving offers the advantage of exact, bit-for-bit copies that can be made with no signal degradation. As masters start making their way into the cloud, it's important that the best possible high resolution copies be archived.

When remastering for iTunes, 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 and finished files on an iOS device.

Although it's possible to remaster from a previously mastered CD source with positive results, in order to qualify as Mastered for iTunes remastered content must begin with a high resolution digitization of the original analog source and must sound noticeably superior to the previously released version. Songs and albums submitted to the iTunes Store as remastered content will be reviewed to ensure that the sound quality shows discernible improvement.

Mastering Tools

Mastering is a series of creative choices. As a professional you can be confident that the work you put in to mixing and mastering your material will be treated just as carefully by Apple. Our role is to faithfully and transparently reproduce the audio you deliver for iTunes—and to make your job easier, you're being provided access to the same world-class tools used to encode the entire catalog for the iTunes Library. Whether you opt to use an automated droplet or a command-line utility, the following tools will help you create the intended sound for your masters:

- **Master for iTunes Droplet.** The Master for iTunes Droplet is a simple, standalone drag-and-drop tool that can be used to quickly and easily encode your masters in iTunes Plus format.
- **afconvert.** The afconvert command-line utility can be used to encode your masters in iTunes Plus format.
- **afclip.** The afclip command-line utility can be used to check any audio file for clipping.
- **AURoundTripAAC Audio Unit.** The AURoundTripAAC Audio Unit can be used to compare an iTunes Plus file to the original source audio file to check for clipping.
- **Audio to WAVE Droplet.** The Audio to WAVE Droplet automates the creation of audio files, in WAVE (Waveform Audio File) format, from any audio file (such as MPEG or CAF files) natively supported on Mac OS X.

The following sections include instructions for using some of these tools to convert audio to iTunes Plus AAC and checking audio for clipping. Both tools utilize afconvert, which is a part of the Core Audio frameworks in Mac OS X. Changes made to these tools are incorporated into Mac OS X by means of Software Update. Software Update is the best way to keep your systems up-to-date so that the system you're using for audio encoding is the same one used by the iTunes Store.

At the time this document was written, Mac OS X 10.6.8 or later is required to have access to the most current version of the AAC encoder.

Master for iTunes Droplet

You can use the Master for iTunes Droplet to automate the creation of iTunes Plus format masters. The Droplet creates an AAC audio file from an AIFF or WAVE source file by first generating a CAF (Core Audio File) rendered with an iTunes sound check profile applied to the file. If the sample rate of the source file is greater than 44.1 kHz, it's downsampled to 44.1 kHz using our mastering-quality SRC. Next, it uses this newly rendered CAF to render a high quality AAC audio file. Once the final AAC audio file is generated, the intermediary CAF is deleted.

To use the Master for iTunes Droplet, drag and drop AIFF or WAVE format source audio files, or folders containing those files, onto the droplet. This droplet works by automating the use of the `afconvert` command-line tool shown below.

For more information about the Master for iTunes Droplet, see the Master for iTunes Droplet Read Me file included in the Master for iTunes Installer.

afconvert

The `afconvert` utility is command-line tool that will let you encode your masters using precisely the same technology used to encode files for the iTunes Store. The `afconvert` utility is built into Mac OS X and can be accessed using the Terminal application.

For more information about `afconvert`, type `afconvert -h` on the command line in the Terminal application on Mac OS X. For more information on the formats that `afconvert` can convert to, type `afconvert -hf`.

Using afconvert to Convert Audio from LPCM to iTunes Plus AAC

Use the following steps to convert your audio file from LPCM to iTunes Plus AAC using the command-line interface. The items in italics are placeholders for filenames. The input file is assumed to be called *source.wav*, and any intermediate files are carried through from the output of one command to be the input of the next. The final output file is called *final.m4a*. These steps assume that your current directory is the same directory as your input file.

If starting with WAV LPCM file at 44.1kHz sample rate:

1. To convert to a .caf file and add Sound Check information, in Terminal, type on one line:

```
afconvert source.wav intermediate.caf -d 0 -f caff --  
soundcheck-generate
```

Or if starting with WAV LPCM file at 48kHz or 96kHz sample rate:

1. To downsample to 44.1kHz LPCM using optimal sample rate conversion and add Sound Check information, in Terminal, type on one line:

```
afconvert source.wav -d LEF32@44100 -f caff --  
soundcheck-generate --src-complexity bats -r 127  
intermediate.caf
```

Then, encode to AAC:

2. To convert to iTunes Plus AAC, type on one line:

```
afconvert intermediate.caf -d aac -f m4af -u pgcm 2 --  
soundcheck-read -b 256000 -q 127 -s 2 final.m4a
```

Previewing the Converted Audio File

The AAC-encoded M4A file (*final.m4a*) created by *afconvert* can be previewed in iTunes or another application that plays M4A files (such as 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 need to decode the AAC-encoded data 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 in iTunes. It should not be submitted to the iTunes Store as an audio source.

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.

Regardless of the bit depth of the original source file (16- or 24-bit), you should generate a 24-bit file to preserve maximum fidelity resulting from the AAC coding process.

Using afconvert to Decode Your AAC-Encoded M4A File

Use the following command to decode your AAC-encoded M4A file. (Note that you can also use the Audio to WAVE Droplet to automate the process of decoding one or more of your AAC-encoded M4A files.) 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.



iTunes won't reject a master file based on the number of clips the file contains. This tool is provided so you can make an informed decision about whether to submit an audio file or go back to the drawing board and make adjustments. It's a creative decision—one that's entirely up to you.

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.
4. 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 original audio file's 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 iTunes Plus AAC against a source audio file. 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. AU Lab is available as a free download at [apple.com/iTunes/Mastered for iTunes](http://apple.com/iTunes/Mastered%20for%20iTunes). For more information on using the AURoundTripAAC Audio Unit, see the AURoundTripAAC ReadMe file included in the Mastered for iTunes Installer.

Audio to WAVE Droplet

The Audio to WAVE Droplet lets you automate the creation of WAVE audio files, in WAVE (Waveform Audio File) format, from any audio file (such as MPEG or CAF files) natively supported on Mac OS X. For example, the Droplet can be used to decode AAC files to 24-bit WAVE files, automating the decoding process used by the Mastered for iTunes 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.

For more information about the Audio to WAVE Droplet, see the Audio to WAVE Droplet Read Me file included in the Master for iTunes Installer.