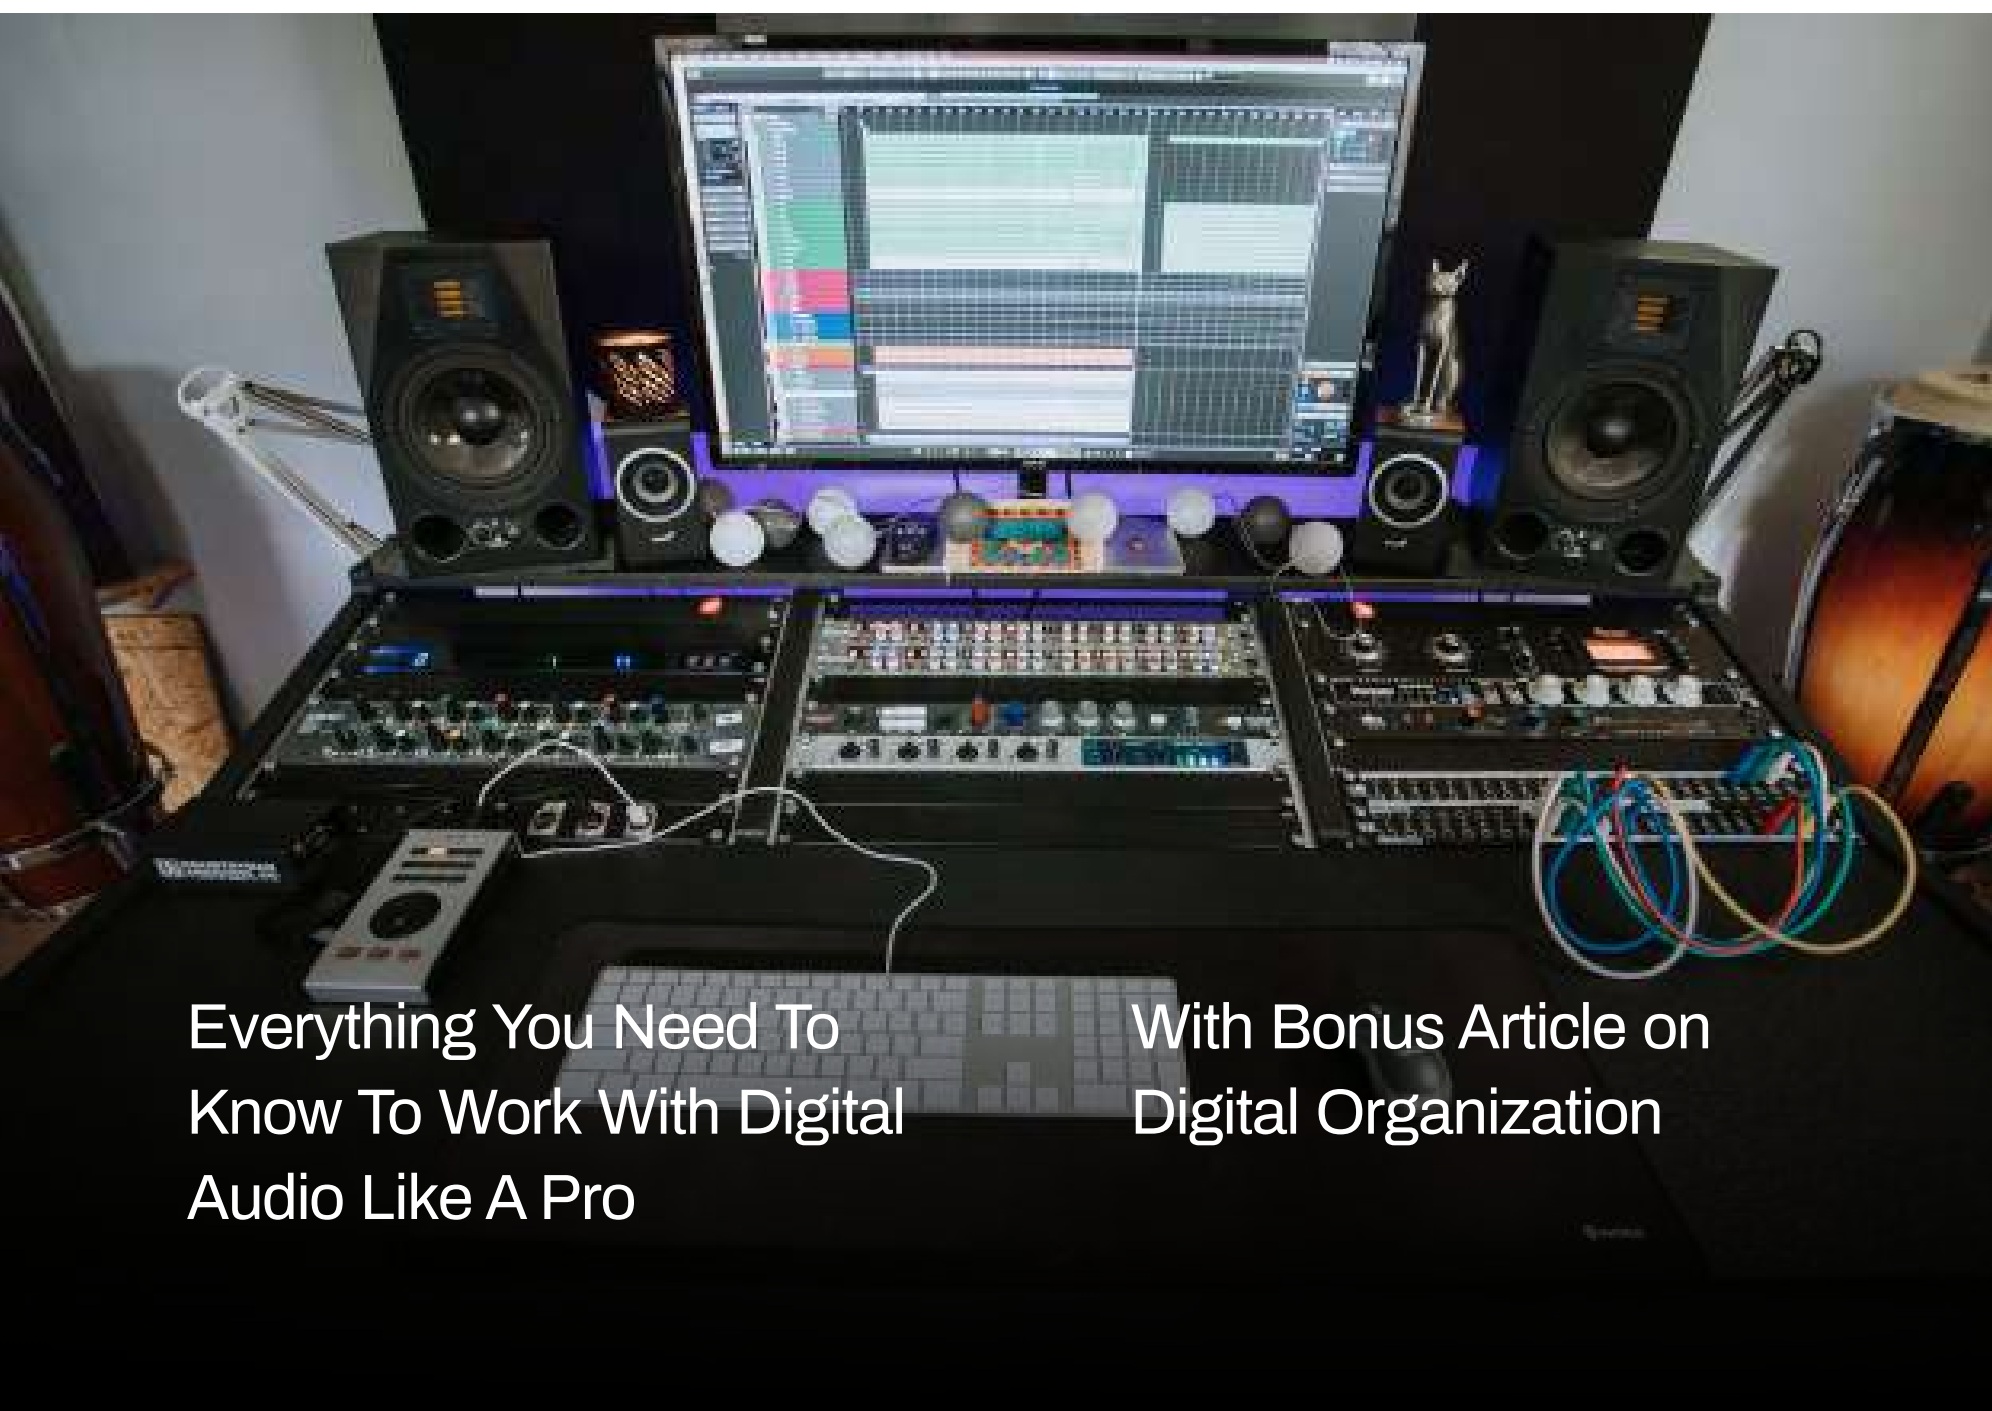


Digital Audio Demystified



Everything You Need To
Know To Work With Digital
Audio Like A Pro

With Bonus Article on
Digital Organization

Table of contents

Chapter 1 – Fundamentals of Digital Audio

Sample Rate 5

Bit Depth 10

Apple Digital Music 17

Chapter 2 – Bonus! Never Lose a File

Backup Strategies 28

Introduction

Humans have created a world defined by bits—tiny perfect measurement of the tangible things around us. These digital bits represent the sounds and sights that we experience in the physical world and over the last few decades our ability to accurately capture reality with these bits has improved dramatically. In music, we can clearly define, in bits, the elements that make up the whole of the sonic picture. We understand how the much-loved analog systems add, smear or distort audio in a way that pleases us and we can create digital effects that come very close to those same analog spices. We must, however, understand how to best capture sound into digital bits so that we can base our recipe on the purest ingredients. Fortunately, there are only a few simple principles that define the digital capture process. What follows are the basics that will optimize your use of digital audio processors. Also included here is a bonus primer on creating a well-organized file and backup system for all your audio projects.


“All things are artificial,
for nature is the art of God.”

Thomas Browne

Chapter 1

The Fundamentals of Digital Audio

Digital audio can be perfect. Analog audio can be perfect. Both of these statements are true, but the meaning of perfect is different in each case. We strive, as an artist for the latter perfect, while as the technician we strive for the former. Without being philosophical, all that means is that we need to first capture the most accurate digital signal we can and then we can “perfect it” however we see fit. Digital audio can be defined by two parameters: bit-depth and sample rate. These parameters, when set properly, provide us with the right amount dynamic range and frequency response to perfectly represent any audio signal. Master these two parameters and the rest is easy.



Understanding Sample Rate

by Adam Kagan

We are all familiar with the two audio file descriptors sample rate and bit depth. While these specifications sound routine, I often get questions from producers and mixers about the optimum settings for a given project. This article will cover the basics and best practices for setting sample rates. Don't worry, though, in another article, we will cover bit depth!

Sample Rate Defined

Sample rate tells us how many times per second we take a measurement of an analog audio waveform as it is converted to a digital signal. Since sample rate has a speed, or frequency, the sample rate defines the frequency response of an audio recording. Specifically, the Nyquist Theorem states that the highest frequency we can record is half of the sampling rate. This means a sample rate of 44.1 kHz can record audio signals up to 22.05 kHz. Accordingly, a 96 kHz sample rate allows for 48 kHz of audio bandwidth. If we attempt to record above half the sample rate, or the Nyquist limit, audible artifacts called aliases occur. Analog to digital converters eliminate aliasing by low pass filtering the analog signal at half the sample rate. This low pass filter is referred to as an anti-aliasing filter. In practice, the low pass filter requires a range to operate, so we state 20 kHz as the practical upper limit for a 44.1 kHz sample rate.

We know that human hearing covers from about 20Hz to 20 kHz, so why would we need sampling rates above 44.1 kHz? One answer is that many people, including scientists, claim that humans can perceive sounds as high as 50 kHz through bone conduction. That claim may theoretically be correct, but through air humans only hear up to about 20 kHz. The second reason is a more practical one. The low pass anti-aliasing filter is not a perfect filter, so it creates some of its own distortions. There is a design tradeoff between how steep a filter can be vs. how little phase distortion the filter produces.

Recommended Sample Rates

Since our hearing is only capable of 20 kHz, can't we just stick to a sample rate of 44.1 kHz? We just learned that sample rates above 44.1 kHz may sound better simply because the analog filter design in the A-D converter has less impact in the audible range. In other words, 44.1 kHz captures all the audio bandwidth humans can hear, but the low pass filter may adversely affect audio below 20 kHz. For this and other reasons, it is recommended that we produce and mix pop music at 48 kHz. First, 48 kHz allows for better sounding anti-aliasing filters than 44.1. Second, 48 kHz uses only slightly more disk space than 44.1. Third, videos usually require 48 kHz audio and much of our audio will be embedded in a YouTube or other video as part of distribution. If you produce music solely for audio CDs, then 44.1 kHz would be the recommended way to go.

Higher Sampling Rates

For audiophile jazz, classical, world music, and some sound design projects, I would recommend the 96 kHz sample rate. This sample rate all but eliminates audible high-frequency aliasing and filter-induced distortions. Furthermore, 96 kHz audio files may provide lower processing latency and allow for excellent sounding downward pitch-shift effects for sound design and game audio. Additionally, a 96 kHz recording can be neatly reduced to 48 kHz when the need arises. If you feel the need to record at sample rates above 96 kHz, you should spend a considerable amount of time testing analog recording chains, converters, and DAWs to find a workflow that suits your purpose. Sample rates above 96 kHz may be more susceptible to jitter problems and will certainly tax your CPU, reduce your track count, and provide fewer plugin choices. I would generally recommend against 176 or 192 kHz unless you have truly studied the pros and cons of those high sampling rates. For reference, the [Grammy's Recording Academy Recommendations for Hi-Resolution Music Production](#) document proposes a minimum sample rate of 48 kHz and a preferred sample rate of 96 kHz for hi-res audio production and delivery.

Sample Rate Conversion

Sometimes sample rate conversion is unavoidable and many software utilities provide excellent sample rate conversion. From a recent survey of mastering engineers, some of the popular sample rate converter programs are [Voxengo r8brain](#), [Weiss Saracon](#), Pro Tools SRC (using Tweak Head settings), [Izotope Resample](#), and [SoX](#). Many other programs provide excellent results and DAWs are continually improving their SRC algorithms.

So we see that choosing a sample rate is relatively straight-forward. Here is a cheat sheet to help keep it all organized:

Sample Rate Cheat Sheet

Recommended sample rates for various situations:

Recording

For pop music stick to 48 kHz, but 44.1 kHz is acceptable. For audiophile music or sound design, you may prefer 96 kHz.

Mixing


Mix sessions should remain at the sample rate of the recording. You will not improve the sound of a project by upsampling a session to a higher sample rate session. If you are mixing on an analog console, print your mix at either 48 kHz or 96 kHz as described above

Mastering

Do not upsample during mastering. For in-the-box mastering, master at the same sample rate as the delivered project. For analog mastering, play out the digital file at its native sample rate, process via analog processors and capture the result at the sample rate the client requires, usually 48 kHz or 44.1 kHz. Also, capture a 96 kHz file for archiving. Do not sample rate convert the final master unless it is unavoidable.

Distribution

Video producers usually require 48 kHz files while digital music distributors like iTunes will accept any sample rate of at least 44.1 kHz. Do not sample rate convert the final master unless it is unavoidable.



“

Every particular in nature, a leaf, a drop, a crystal, a moment of time is related to the whole, and partakes of the perfection of the whole.

Ralph Waldo Emerson

Understanding Bit Depth?

by Adam Kagan

In the previous article, we discussed how sample rate and frequency response relate to each other. This article will discuss how bit depth relates to another important quality of audio—dynamic range. Dynamic range is just a fancy term for the difference between the loudest sounds and the quietest sounds we can record. Read on to learn how to optimize bit depth.

Bit Depth Defined

As already noted, bit depth defines the dynamic range of an audio signal. Specifically, each bit represents approximately 6dB of dynamic range, so 16-bits provide 96dB of dynamic range, and 24-bits provides 144dB of dynamic range. Bits, or binary digits, are based on exponential math, and a 16-bit file can measure 65,536 values (2^{16}), while a 24-bit word measures over 16.7 million values (2^{24}).

In the real world, professional analog recording and mixing devices have a typical measured noise floor somewhere between -120dB and -90dB, so as far as recording analog audio, we don't take full advantage of the 144dB of dynamic range that 24-bit provides. However, once audio is digitized into production software where we process and mix signals together, the additional dynamic range of 24-bit or higher mathematical operations becomes very important. In other words, 16-bit recordings can capture all the dynamic range of typical analog audio signals, but more bits are needed once we start processing the audio inside a DAW.

I receive many files from clients who record and produce multitrack sessions at 16-bit resolution. When I mix or master 16-bit files, I find that creating a new session at 24-bit or 32-bit floating-point and importing the 16-bit audio files sounds much better than working inside a 16-bit session. As a rule, I mix and master every project in a 24-bit or 32-bit floating-point session. Pro Tools, Logic and other DAWs usually process audio with 32-bit floating-point (or higher) precision. My advice is to start your projects at 24-bits and you won't have to convert later. Sure, 24-bit files use more disk space than 16-bit, but that shouldn't be a consideration with modern disk sizes.

What is Dither?

We can't discuss bit depth without mentioning dither. Capturing an audio waveform as a digital signal is a process that converts a continuous (analog) waveform into many individual (digital) measurements. These discrete measurements don't measure every point along the smooth analog waveform and therefore create a distortion called quantization errors. We can randomize these errors, effectively making the distortion undetectable by adding a very small amount of random noise to our digital signal. This noise is called dither and is typically only 3dB above the noise floor, or about -93dB for 16-bit audio. What we wind up with after dithering is a perfect representation of the analog waveform plus some very low-level noise. It would seem like we can't recreate a perfect analog signal if we only measure the signal at the digital sample rate, but the Nyquist Theorem (described in our "Understanding Sample Rate" article) states that an analog signal that has been digitally sampled can be perfectly reconstructed from the samples.

DAWs process audio at 32-bits (or higher) precision, so if we need to export a mix, master, or stems from our DAW at less than 32-bit resolution, we need to apply dither to eliminate quantization errors. As a rule, always apply dither when exporting a file from a DAW to a lower bit depth than your DAW's native resolution. If you are exporting audio as a 32-bit float file from your DAW do not apply dither. If you are freezing tracks or using bounce-in-place to hand off files to another engineer, you must look into how your specific DAW handles those files.

As a rule, always apply dither when exporting a file from a DAW to a lower bit depth than your DAW's native resolution

Can We Hear Dither?

Mastering engineers agree that properly applying dither improves the quality of audio that is reduced from higher bit depths to 16 bits.

Dither noise is random noise and effectively 93dB below your peak levels, so only in extreme circumstances will the actual dither noise be audible. However, if you did have a case where dither noise is somehow audible, remember dither is random noise, and humans are excellent at ignoring random noise.

Alternatively, if we truncate a file without adding dither, we will create distortion that is musically related to the audio file. Distortion that is related to our audio is not like random noise, but instead will reduce the quality of our audio, usually in an ugly way.

There are many flavors of dither, including different types of noise shaping (EQ'd noise), and if you are so inclined, you should perform your own listening tests to see if different types of dither create audibly different results. In my experience, I can sometimes tell that a file was exported without dither, but I have never been able to tell the difference between various brands or flavors of dither that is properly applied. In the end, applying dither is as easy as pressing a button, so apply it when needed and don't give it a second thought.

Conclusion

Choosing a bit depth is relatively straightforward and the rules about applying dither are also fairly simple. Here is a cheat sheet to help keep it all organized:

Bit Depth Cheat Sheet

Recommendations for various situations:

Recording

For close-miked and amplified music recording 16-bit may be acceptable, but 24-bit is recommended. For live to 2-track or extremely dynamic recordings always use 24-bit.

Mixing

Mix sessions should be set to at least 24-bits of resolution. Most DAWs process internally at bit rates of 32-bit floating-point or higher.

Mastering

Deliver 24-bit or 32-bit floating-point mixes to the mastering engineer and the mastering engineer should work at the highest bit depth possible.

Dithering

Always apply dither when exporting a file from a DAW at a bit rate lower than the DAW's bit rate. Do not apply dither when exporting a 32-bit floating-point file

Distribution

16-bit files are acceptable for PCM audio formats, like audio CDs. If a file will be converted to a lossy format, it is best to start with a 24-bit or 32-bit file before encoding. Check with your particular distribution service or aggregator to see what formats they will accept.

Archiving

NARAS recommends archiving PCM audio files at the original sample rate and highest bit depth possible, like 44.1 kHz, 32-bit floating-point.



“

Digital for storage
and quickness.
Analog for fatness
and warmth

Adrian Belew

Apple Digital Masters

by Brad Pack

Editor's note: As of August 2019 Apple has rebranded its iTunes music app and updated the Mastered for iTunes content specifications and this document reflects those updates. We appreciate any and all feedback and strive to provide useful, accurate and up-to-date information for our readers. Apple has not discontinued iTunes, but merely created a home, called Apple Music, for its streaming services and store for music, podcast and audiobook content.

Mastering is a complicated process — a delicate blend of art and science that requires a well-trained set of ears, a deep technical understanding of digital and analog audio, and, perhaps most importantly, the ability to assess the artistic intent of a piece of music.

The mastering engineer's job is to shape the sonic presentation of a finished mix and optimize playback across all systems and media formats. The subsequent mastered version of the song acts as the master copy from which all copies or duplicates are derived—files for online streaming, CDs, and maybe even vinyl.

Each distribution platform, analog or digital, has its own specifications and standards regarding file type, overall loudness, and included metadata. One of the most stringent sets of requirements, or “best practice” standards, comes from Apple via their Apple Digital Master specifications. In this blog, we'll take a look at the requirements for earning the Apple Digital Master badge, what those requirements mean, and how to make sure you're delivering the highest quality digital audio files possible. You can safely assume that if your masters meet the Apple Digital Master requirements, the masters will also be of superior quality for any other streaming or online distribution platform.

What is an Apple Digital Master?

[Apple Digital Masters](#), formerly referred to as Mastered for iTunes (MFiT), refers to a set of guidelines developed by Apple in order for mastering engineers to create high-quality masters for streaming on the Apple Music service and for sale through the iTunes Music Store. Apple has recently announced a change to the iTunes store, but don't fear—the iTunes music store is not going away, but will reside inside the new Apple Music app and the video content will be available via the new Apple TV app. The idea is that the masters we provide to Apple can be converted to Apple's AAC format while retaining the original dynamic range and fidelity that exceeds 16bit 44.1kHz red book CD quality standard.

Apple uses an audio format known as Advanced Audio Coding (AAC) for compressing and encoding digital audio. While similar to mp3, AAC files offer improved encoding algorithms over MP3 files and typically provide better audio quality for similar-sized encoded files. In 2003, when the iTunes Store was first launched, Apple used 128 kbps AAC files, but over time the iTunes catalog was upgraded to iTunes Plus, which uses a 256 kbps (VBR) AAC encoding format.

AAC and MP3 descriptions

The term **kbps**, or kilobits per second, describes the data rate of an mp3 or AAC, and therefore also describes the size of the compressed audio file. A 16-bit 44.1 kHz uncompressed .wav (referred to Linear PCM) file has a data rate of 1.4 megabytes per minute, so a 3-minute song would have to stream just over 30 megabytes of data. The standard 256 kbps AAC file streams only about one-fifth of that data for the same song. Common mp3 bit rates include 128, 192, 256 and 320 kbps, while AAC is optimized for 256 kbps.

The description **VBR** refers to variable bit rate, as opposed to **CBR**, or constant bit rate. Variable bit rate is simply an efficient way to create smaller files without sacrificing audio quality. The bit rate of a VBR file varies as a song streams to supply the required amount of data. When the audio is very simple, a lower bit-rate can be used and when the music is more complex, a higher bit rate is used. AAC uses a type of VBR encoding referred to as Average Bit Rate.

256 kbps AAC audio files typically sound as good or better than larger 320 kbps mp3 files. AAC files also support multi-channel audio formats, as well as wider frequency response and superior transient response, as compared to mp3 files.

Sample Rate

The sample rate of an audio file describes the frequency response of the captured audio, and we want to reproduce frequencies up to at least 20 kHz. To reproduce 20 kHz, the sampling rate must be at least double the frequency, so a minimum sampling rate of 44.1 kHz must be used. Apple's latest encoding methodology uses sample rate conversion (SRC) to resample your master file to a sample rate of 44.1kHz, regardless of the original sample rate of your supplied master .wav file. Apple uses an excellent algorithm to resample the audio in order to produce the best sounding 44.1kHz file.

Bit Depth

The bit depth of an audio file determines the dynamic range of an audio signal. Each bit represents approximately 6 dB of dynamic range. The most common bit depths for audio files are 16-bit (CD quality), which provides 96 dB of dynamic range, and 24-bit (typical DAW file), which provides 144 dB of dynamic range. Regardless of the audio files bit depth, 32-bit floating-point processing has become common inside DAWs. Without getting into the pros and cons of 24-bit integer math vs 32-bit floating-point binary math, we can safely say that 32-bit floating-point internal processing provides more accurate (better sounding) processing of audio inside a computer. Audio interfaces and analog to digital converters on the market today provide either 16-bit or 24-bit audio to our DAW, but the math inside our DAW may be done at the higher resolution that 32-bit floating-point math provides. When outputting a 24-bit final master from a DAW running at 32-bit float, be sure to dither from 32-bit to 24-bit. Masters may be output from many DAWs at 32-bit float resolution, but be sure that 32-bit audio files can be accepted by the distributor or end user.

How Sample Rate and Bit Depth fit into Apple Digital Masters

In order for Apple to create the best-sounding streaming files, they use a two-step process to convert supplied masters. According to Apple, the file is first sample rate converted at 32-bit float resolution and then the file is encoded into the AAC format. This process avoids clipping and retains the dynamic range of the original file. Additionally, this process avoids the necessity of adding dither to your master.

In order to ensure that your music may be encoded to AAC files as transparently and faithfully as possible, Apple has created a [document](#) that outlines the best practices for delivering files to Apple. Apple Digital Masters certified mastering engineers and providers can be trusted to follow Apple's best practice recommendations, and Apple relies on the honor system to enforce its specifications. Below are the Apple Digital Master delivery recommendations:

- Provide high-resolution masters, ideally at 24-bit, 96 kHz. It's important that you maintain the highest possible resolution throughout the production process, but upsampling files to a higher sample rate won't add information or improve the sound of your project. This means if you create your masters at 24-bit, 44.1 kHz, you should not upsample to 96 kHz. Simply keep the files at 24-bit, 44.1 kHz.
- Avoid clipping at all costs. Make sure your tracks have sufficient headroom. Although Apple will not reject files for a specific number of clips, tracks with audible clipping do not qualify for Apple Digital Master status. Apple suggests leaving 1 dB of headroom below 0 DBFS as your ceiling when creating your final, mastered, .wav file. When using a digital limiter, simply set the ceiling to -1.0 dB (True Peak, or ISP enabled).

- Check your masters on the devices your audience will be using. Take into account the limitations of these devices, as well as the listening environment of your audience. There are many software tools to audition your master as an encoded AAC file, such as Ozone 8's codec preview or Sonnox's Codec Toolbox, among others.

Mastering Tools

In order to make sure your track qualifies for Apple Digital Master status, Apple has provided access to the [actual tools](#) they use to verify the integrity of your master.

The Master for iTunes Droplet is a standalone, drag-and-drop tool for encoding masters in Apple's AAC format. It creates an AAC audio file from an AIFF or WAVE source file. If the sample rate of your .wav file is higher than 44.1 kHz, it's downsampled to 44.1 using Apple's optimized SRC. Alternatively, if you are familiar with using Terminal on Mac OS, creating AAC files can be done manually using the command-line utility `afconvert`. Examples of these command-line codes are included in the [Apple Digital Masters document](#).

In order to make sure your track qualifies for Apple Digital Master status, Apple has provided access to the actual tools they use to verify the integrity of your master.

`Afclip` is another command-line utility that can be used to check files for clipping. `Afclip` generates a unique stereo audio file that uses the left channel for the original audio and the right channel for a graphic representation of each clipped sample so you can quickly identify when and where a file is clipping. It also produces a text readout with each instance of clipping, as well as a summary of how many total clipped samples the audio file contains.

`AURoundTripAAC` is an Audio Unit plug-in that lets you compare audio encoded using iTunes Plus AAC against its source file to

quickly A/B the changes. It can be used in any audio unit host application, such as Logic or AU Lab.

The Audio to WAVE Droplet is a standalone tool that creates .wav files from any audio file natively supported on Mac OS X. It can be used to decode AAC files to 24-bit .wav files just like afconvert, and allows you to decode multiple files at once.

Future-Proof Mastering

Delivery requirements for mastering engineers are constantly evolving. As technology advances and better sounding formats are introduced, mastering engineers must adapt to the new standards. Back when CDs were the de-facto format, mastering engineers were only required to submit 16-bit 44.1 kHz files. Today, every distribution service and aggregator has its own set of standards for submitting digital audio files.

As a mastering engineer, it's typically considered best practice to provide multiple files for clients. While the 16-bit 44.1 kHz file may be best for CD duplication, it is not a high-resolution format and is not best for conversion to compressed streaming formats like AAC or mp3. Along with a properly dithered 16-bit file, you should provide your client with an additional high-resolution Apple Digital Master quality file to make sure your work will sound great on high fidelity playback systems, as well as when encoded to lossy distribution formats.

The Apple Digital Master Badge

Apple's Digital Master rules are simply best-practice guidelines. As previously mentioned, we are on the honor system to produce and deliver the best quality masters possible. Apple provides a list of Apple Digital Master providers to their content developers, but any facility can produce Apple Digital Masters as long as they understand and abide by Apple's guidelines. Unfortunately, if you would like the iTunes store to display the "Apple Digital Master" badge for your music, things get a little complicated. Albums distributed by major labels will almost certainly display the badge, but if you use an aggregator like Tunecore, Distrokid or many others, you must do your homework to see how they handle your files and if they even provide the option for Apple Digital Master certification. Distrokid, for example, requires that you email them to ask for specific instructions for uploading your file and receiving the Apple Digital Master badge.

In the end, we all desire to distribute the highest-quality version of our music possible. As encoding technologies improve and streaming bandwidth increases, the customer will receive better and better-sounding files, and as long as we follow a best practice, like the Apple Digital Master standards, our master files will always be optimized for creating the best sounding product.

Chapter 1

Key Takeaways

- ☐ Sample rate defines the available frequency response.
- ☐ Higher sample rates may suit more delicate music.
- ☐ Stick to 48kHz sample rate for video projects.
- ☐ Bit depth defines the available dynamic range
- ☐ Record, mix, and master at 24-bit or higher
- ☐ Use dither to output a 16-bit master
- ☐ Apple Digital Masters provides best practices for creating HD audio masters
- ☐ Maintain dynamics and headroom for the best sounding distribution copies

A close-up, slightly blurred photograph of A.R. Rahman. He is shown from the chest up, wearing a blue shirt. He is looking down and to the right, with his mouth open as if speaking. A silver microphone is visible in the lower right foreground. The background is dark and out of focus, with a bright circular light source visible in the upper left.

“

“Everybody around is so good, so to create music that will connect with so many people is not humanly possible without inspiration.”

A.R. Rahman

Chapter 2 - Bonus!

Protecting our Work

As musicians, writers, composers, producers, and engineers, we typically run our own business and provide services to others. Our job involves creating and manipulating our own content, as well as that of others. Our clients and fans trust us and rely on us. So, along with providing our creative services, we have also become obliged to create, deliver and maintain digital (and maybe analog) files of our work. By default, we have become technicians and librarians. The value of our work, whether it is a symphonic score or one line of voiceover, is more valuable than simply the time spent or the hard drive space. The engineer who can locate and deliver a revised master from a 5-year old project is a hero. The engineer who can't may lose their client's trust forever. Maintaining and backing up our computer systems and files has become an obligation as well as a nuisance, and it behooves us all to become fluent, habitual users of some sort of backup and storage routines. The following article should help shed some light on the process that one engineer, Tiki Horea, uses to maintain his sanity amid the world of his digital files.

Backup: Strategies for Never Losing Files

by Tiki Horea

A proper backup strategy is something that distinguishes pros from amateurs—or at least the wise from the unwise! Don't you hate it when you lose days/months/years (!!!) worth of work because your hard drive died and your files are gone? You share your pain on Facebook groups and while the first few replies are sympathetic, you quickly get smacked over the head with the truth: "You should've backed up."

In this article, you'll learn how to save yourself time, frustration, and money by never losing files. Even more importantly, you won't lose your client's files or their trust.

Story #1

I remember when I moved from Windows to macOS and I wasn't used to how things run in the new environment. I had backups of all my projects on an external hard drive, and that was my whole backup strategy.

At one point, I wanted to restore a few projects from my external backup drive back to my main drive. In the past, when Windows asked whether I want to merge the folders and their contents or not, I simply clicked "Yes" and all went well. When my new macOS machine asked me if I want to replace the files, I simply agreed to it.

My old Windows habit of clicking “Yes” lead me to delete every song I had ever written. ALL OF THEM.

After spending some time sulking that I had lost my songs, I moved on. Hey, things happen. The time spent sulking meant not writing for another year and a half or so. Good thing nowadays I’m a mixing and mastering guy, eh?

So, don’t you put off sorting out your backup situation as I did. Procrastinating is always a bad idea. There are two more disaster stories at the end of this article, and if this one didn’t do it, they will most likely convince you to save your future self a lot of time, frustration and money.

Be your future self’s best friend, will ya?

Keep Your Drives Organized

The way you arrange your files in your system is very important for your peace of mind and overall efficiency. Proper file management and keeping track of where every file lives will end up saving you a lot of time over the course of a year. For instance, a simple organization strategy would be like this: Create two parent folders and name them “In Progress” and “Finished Projects.” Each time you begin a new project, simply create a subfolder in the “In Progress” folder named for the artist, client or project. After you finish the project, move that artist folder to the “Finished Projects” folder. This method makes finding particular projects very simple. Naming files and folders is also an important consideration and we will get to that in a bit. For my music projects, I have a dedicated “Work Drive” that houses three parent folders:

Folder 1: A_Work_In_Progress

This folder contains subfolders such as:

- Mastering
- Mixing
- Practice
- Production
- Recording

Folder 2: Finished

This folder, unsurprisingly, is where I archive finished projects. This folder contains the same subfolders as *A_Work_In_Progress*

At the end of every project I create some files to be archived and possibly delivered to the client:

- 1 Bounced/rendered/exported stems
- 2 All tracks in the session are frozen/committed
- 3 All related project folders are moved into the main project folder
- 4 The project folder is compressed so that no audio files get misplaced.

Folder 3: On_Hold

This folder has random projects that are on hold for one reason or another. Self-explanatory, I know.

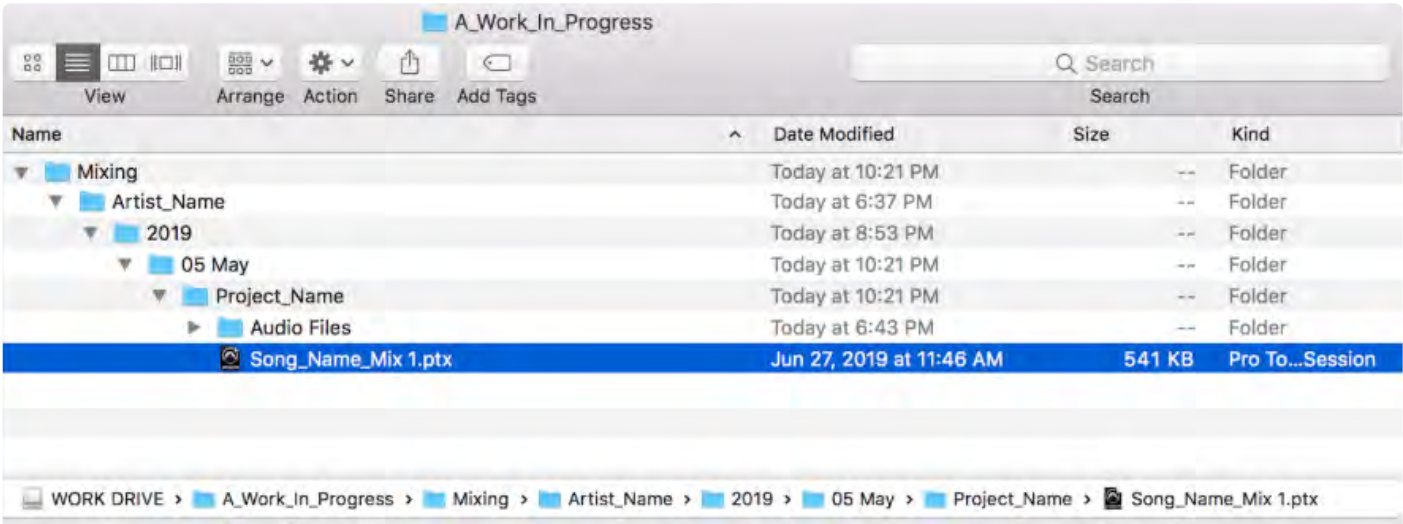
Folder Hierarchy

In my case, each of the three parent folders contains artist folders labeled with the artist or contact person’s name. Inside each artist folder, I have subfolders labeled for the year I received the project. Each year folder contains folders for each month in which I had a project. I number the month folders sequentially to make it easier to find them. Below is an example of my folder structure (Don’t judge me, one can dream):

BAND NAME > YEAR > MONTH > PROJECT NAME > SESSION

Phil Collins > 2019 > 05 May > Album 9 > Song_1

Example folder structure:



File Naming Conventions

Consistency is an important element of organization, so I stay consistent with my folder structure and how all my files are named. To that end, I use [NameChanger](#) for my batch renaming needs within macOS, and [AdvancedRenamer](#) when working in Windows. Here are some of my batch renaming actions:

- Remove the “_01” suffix from files exported from Logic Pro or Pro Tools. My_Song_final_01.wav becomes My_Song_final.wav
- Remove number prefixes from the audio files, like 01_Kick, 07_Tom5, etc.
- Conserve characters by using reasonable abbreviations. I replace the word “guitar” with “gtr”, as in El gtr, Ac gtr, Gtr clean, etc. Short file and track names allow the names to be more easily displayed on crowded screens.

How to properly backup: The 3-2-1 Rule

Follow the 3-2-1 backup rule. The 3-2-1 backup rule is an acronym for a sure-fire strategy to keeping your data safely backed up.

- Keep at least **3** current copies of your data
- Store **2** copies on different storage media
- **1** of the copies must be stored offsite

3 copies of your data means that you have your primary data and two backup copies. Remember that all storage devices fail at some point and using different types of storage media decreases the likelihood that multiple backup copies will fail at the same time. Therefore, the two backup copies should be different storage types, like external SSD or cloud storage.

Disasters such as fire, theft, and floods can destroy all local backups. This means you must store one of your backups offsite. Cloud storage subscription with companies like Backblaze, Carbonite and IDrive provide secure, offsite storage plans with automated backup options. Alternatively, your offsite backup can simply be a drive that you take home with you each night. If you work at home, you could keep that drive at a family member's home or in a secure storage area in a separate building, like your garage.

“

“If your payloads cost hundreds of millions of dollars, they actually cost more than the launch... Reliability becomes much more important than the cost.”

Jeff Bezos



Putting It All Together

For audio professionals, your backup strategy could look something like this:

- Your audio/session drive should be a separate drive from your system drive—don't cheat on this, even if you work on a laptop.
- Your audio/session drive should be synced to a cloud-based backup service that is constantly backing up your data.
- Create a backup schedule with Carbon Copy Cloner or ChronoSync to clone your work drive to a backup drive, which is attached to your system or your local network.

Here is how I set up my system and backup workflow:

Backup 1

[CarbonCopyCloner](#) backs up the entire A_Work_In_Progress folder, hourly, to my local Google Drive folder. As you might expect, my Google Drive folder automatically syncs and uploads to Google's cloud servers.

Backup 2

[CarbonCopyCloner](#) also clones A_Work_In_Progress to my "Work Backup" drive.

Backup 3

[Backblaze](#) online backup is set to continually backup A_Work_In_Progress folder to a secure cloud storage server.

While the 3-2-1 backup rule is not 100% fail-safe, it is the best practice recommended by US government authorities such as the [United States Computer Emergency Readiness Team](#). IT security experts also recommend this strategy as the best way to keep your data safe from disappearing off the face of the Earth.

Manual vs. Automatic Backups

You could backup manually, but this means you must remember to do it. Automating the backup process with a software solution ensures consistent backups and reduces the stress of having to remember to backup. I use [CarbonCopyCloner](#) on the Mac to automate all my backing up. [Acronis True Image](#) looks like the comparable software for Windows.

Keep Older Versions of Your Files

It's very important that your backup solution maintains multiple versions of your files; meaning older versions are kept, even when newer versions are backed up. What if you accidentally delete an important file before you backup or simply wish to go back to an older version of your project? Versioning to the rescue!

System Drive Backup

Backing up your project data is vital; however, do not neglect to backup your system drive. Backup programs like [CarbonCopyCloner](#) and [Chronosync](#) can create a bootable clone of your system drive. A bootable clone can be a godsend if your system drive fails or becomes corrupt. If you have a bootable backup, you can simply swap your boot drive, restart, and you're back to work.

For both versioning and system restore, one could use Apple's built-in Time Machine utility. For things like restoring previous copies of a file, I find Time Machine easy to use, surpassing both [Backblaze](#) and [CarbonCopyCloner](#). However, Time Machine will not create a bootable backup of your system drive and is not the most efficient solution for backing up audio drives.

Conclusion

- Spend the time to develop an organized drive and folder structure.
- (Re)name your files so that they work for you, and not against you.
- Most important of all, is to backup properly, using the 3-2-1 rule.
- Do not wait until disaster strikes before you take this seriously.

Retrieving old sessions for a client not only adds value to your services but also earns you some brownie points. In fact, you'll earn brownie points even if you charge for the service!

Please backup. Now. Be your future self's best friend, will ya?

Story #2

Once upon a time, there was this super awesome music producer called Patch. He was asked to record a few songs for the incredibly talented Nicola Brown so he called upon some of his most trusted engineer friends, Kevan and Tiki. Together they recorded a beautiful sounding drum kit – one of the best sounds ever captured by anyone, ever.

Patch, being the lovely, caring person he is, decided to collect the session files the following day as it was late and everyone wanted to go home. Alas, on his return he found no trace of the Pro Tools project recorded the previous evening. Gone were the 6 hours worth of audio. Gone was the best drum sound ever. Gone was some nice sounding acoustic and electric guitars. Gone was Patch's happiness.

Close to tears, he booked the studio for the following week and re-tracked the whole EP. It sounded much better but was still a huge headache. So remember kids, always backup your sessions, no matter what. – [Patch Boshell](#)

Story #3


A few months ago I got an external hard drive specifically to transfer work between my office computer and my laptop. At first, I was copying and pasting the work between the computers and the external hard drive, effectively creating 3 backups. As time went by I got lazy and started working directly from the external hard drive instead.

You can predict what happened next! External hard drives are fragile. One fateful day I dropped it. When I plugged the drive into the computer it refused to mount at all! Imagine how furious I was with myself. I had to recreate a bunch of client work and wasted too much precious time...and money. – [Matteo Depares](#)

Chapter 2 - Bonus Content

Key Takeaways

- ☐ Organize your drives and projects into clear folders
- ☐ Consistently name your project files in a descriptive manner
- ☐ Follow the 3-2-1 backup strategy
- ☐ Rely on automatic backups
- ☐ Ensure your system is backed up as well as your media files



“

“Music is powered by ideas. If you don't have clarity of ideas, you're just communicating sheer sound.”

Yo-Yo Ma



If you're looking to get the best out of your studio, then try this...

The same know-how we've used to make this e-book can be used to vastly improve your monitoring. Our Reference 4 calibration software makes your headphones and monitors coloration free, so you can hear what's really happening in your mix.

Sonarworks Reference 4 uses acoustic measurements and applies them to your headphones and monitors, so any tonal coloration is reversed. The result is a clear picture of your work, so you can focus on doing what you love!

Available in major DAW plug-in formats and as a standalone Systemwide app Reference 4 is already used in more than 35 000 studios worldwide.

Click to Try Reference 4 for Free

or visit sonarworks.com/reference