

# Digital Audio Levels

v2.1  
July 28th, 2014

I have to admit I struggled a bit when trying to decide if I should make this guide a Part 2 to my original guide on leveling and meters, or if I should just start from scratch with a new revision. The first version was written mainly to help digital DJs with laptops understand digital audio better, along with some related topics along the way. However, in my attempt to keep things basic and easy to understand, I perhaps simplified things too much (and too poorly), which lead to a lot of follow up questions. Also, it was apparent more producers were reading the guide than DJs, and that some of the concepts could have been more user-specific. So, in the end it seems more appropriate to just start from scratch and erase the mistakes of the past.

I'll be splitting this guide up into two main sections, one focusing on performers using digital audio system to DJ or play live, and one focused mainly on producing and digital in the studio. The main difference in my opinion is that performers are most concerned about audio coming OUT of their systems, and producers are more concerned with getting audio INTO their DAWs. Certainly there's going to be parts of both sections that apply to both groups, so it's probably worth reading them both no matter what your background. First up then, the performance section...

## Digital levels and meters on stage and in the DJ booth

I'll cut right to the chase on this one, the number one thing most live performers want help with is making their laptop-based performances as loud as those people using turntables or CD decks. It's understandable, you show up at the club, plug in your gear, and you can't help but notice that you're going to be quieter than the DJ who was playing before you. It will be much easier to understand what is going on if you accept two facts right away:

- A CD deck or turntable is a device with a known and consistent input range, a laptop and soundcard is not.
- A CD deck or turntable is a device with a known and fixed output level, a laptop and soundcard is not.



Let me explain these in a little more detail, so you can understand what I mean. My first point is that when you are using decks (CD or vinyl based) to DJ or perform, the music you play on them will largely be the same volume already. Obviously there's some small differences, but most music these days is mastered and sold in a relatively narrow volume range. The CDs, records, and mp3's you download will all be roughly the same volume, so the manufacturers of decks know pretty well what the maximum level of signals they will be dealing with are, and can design the gear to use this as best possible.

A digital audio workstation (like Ableton Live) used for live performance is a different beast though. It's designed to be useful with a much larger range of volumes, from the quietest recorded whisper, to the loudest recorded explosion. There's also the issue of internal summing in a tool like this, where you might be playing multiple audio tracks at once and need to adjust their volumes appropriately to avoid clipping (more on this later). In short, unless you are always playing store bought CDs and MP3's one at a time, you do not have the same guaranteed input level as a CD deck. This puts the burden of volume and loudness control on the user, and many do not understand the best way to set up their software and prepare their files to handle this in the most appropriate way.

My second point is similar, in that a turntable or CD deck does not have a volume control, the output level is fixed. So for years manufacturers of DJ and PA mixers have been able to build their products ideally suited for this fixed output level they will be receiving. When you know the volume of a signal you will be receiving, you can design the rest of the system around that.


A laptop with a soundcard is of course a totally different story. Most soundcards have multiple ways to affect the output volume, either on a per output basis, or globally. Some even offer different standards which the outputs are calibrated to (you might see this labelled -10 and +4 in your soundcard). Further complicating things, you also have a lot of control of the signal level before it even reaches your soundcard, either with your DAW master fader, channel faders, and even the ability to permanently alter the volume of the digital audio file itself. Setting up everything to provide the cleanest and loudest signal possible at each step of the process is called gain-staging, and we're going to look at the proper way for DJs and live performers to approach this.

The first thing to look at is the audio files you're going to be playing. Like I mentioned earlier, if you're playing songs from a CD or that you bought online, these will already have been mastered and should be close to the right starting level (and when I say level, I mean volume/loudness). For digital vinyl systems like Serato, Traktor, and similar, each of these tracks will likely be sent to its own output on your soundcard, so there's not much more we need to do with them.

But what if you're using Ableton Live, or one of the other programs, to mix internally on your laptop and are only connecting to a PA with a single stereo feed? You will need to make sure that when you play back two more tracks at once, they are not 'clipping' the output of the software before it goes to your soundcard.

## CLIPPING!

I'm going to pause here briefly to talk about clipping, because it's something that is vital to understand when talking about digital audio. Clipping occurs when a signal level exceeds the maximum permissible amount, and distortion occurs. In the analog world, a little bit of clipping is not always a bad thing, it's often much subtler and can even add harmonics to the signal that some people find pleasing. In fact there are quite a few analog boxes costing thousands of dollars that do nothing but add tiny bits of this pleasing distortion. There's a little more leeway in how diligent one needs to be when setting and working with audio levels in the analog world.




However, when you have a digital signal that is clipping, it is being permanently and irreparably distorted and damaged. We're no longer making a perfect copy of the signal we're trying to reproduce, we're distorting it in a very unpleasing way. So as a general rule of thumb, you NEVER want to see clipping in a digital system. That includes not only your software, but also your soundcard.

So, clipping is bad, we all understand that, right?

To avoid this in your performance DAW, you need to play back as many tracks at once that you will realistically play live, and adjust their individual levels until they no longer clip the master track when summed. There's a few ways to do this:

- Lower the volume of the audio file itself, destructively or with a "Clip Volume" control like in Ableton Live.
  - Lower the channel faders for each track.
  - Use a plug-in such as Ableton's Utility on each track to lower the volumes of the audio files before they get sent to the master track or output.
- I always advocate the first method, using the Clip Volume control found in Ableton Live's Clip View (set to -4dB below)



The basic premise is that you want to lower the volume of all of the audio files that will play back at the same time until the master channel meter is no longer clipping. To make things easier later on, once you have determined what this level will be, you might want to go and set all off your tracks to the same volume too. This way no matter which combination of tracks you play, you know in advance that clipping will not be an issue. So, how do we do this?

Well, you can use the channel faders in Live for this, but I personally find a dedicated metering plug in to be easier to use and more accurate.

Looking at the upper right area of the plug-in, you can see two readings for the left and right RMS values. RMS means Root Mean Squared, and is a means for determining average loudness. Two files with the same RMS reading will sound roughly the same volume. So, let's assume we want to set the levels for DJing with only two tracks playing at once. Place one of these meters on each track with an audio file you want to play in your live set, and start playback. Lower the Clip Volumes of the files on each track so that when they are both playing back with the same RMS reading on FreeG, the master meter is not clipping. You should aim for peaks on the master channel about -3dB or so, just to give yourself a little bit of a safety net to prevent clipping.

Also, make sure that the channel faders on each track are set where you want them for your performance too. Some people leave them at 0dB, some people prefer putting them all the way up to +6dB so they can slam a fader up for full volume. It's up to you, just put them where you will want them maxed out when you play live while you do this.

To recap, set the Clip Volumes of both audio files so that they have the same RMS reading on FreeG, while peaking at around -3dB on Live's master meter. Don't forget you can drag the master track header left and right to see the numeric readings of Live's master meter by making the track wider. You can also drag the line above the master meter to make all the meters taller with finer resolutions too. Remember to click the Clip Save button once you're done too, or else your Clip Volume settings will not be remembered later.

Once you have both tracks at the same RMS reading as mentioned above, you can now use the FreeG plug in to measure the RMS readings of all your tracks, and adjust their Clip Volume levels so they are all the same RMS reading as above. This way, no matter which track you want to play later on, you don't have to worry about clipping your DAW, or if they are louder or quieter than each other. I tend to do this process while warping my tracks initially, just to save time. Keep in mind that the RMS readings are only rough guide too, if one track sounds louder than another with the same RMS reading, don't be afraid to trust your ears and compensate some with the Clip Volume until they SOUND the same volume. This is music after all.

So, now we are sure that we are working with a consistent audio volume in our laptop, which is the input side of things I mentioned at the start of this guide. What about the output side though, the part most people are really interested? I want to reiterate that the output of a soundcard or laptop is not supposed to be the same level as the output of a CD deck or a turntable. In fact it was specifically designed to NOT be the same output level. Many people think there is something wrong with their soundcard or set up because it is quieter than a traditional DJ set up, this is not the case. But, that doesn't mean there's not some steps we can take to level the playing field!

I'm going to cover the options you can use to make your laptop-based set up as loud as a traditional DJ set up, starting with the things you should try first:

- Make sure your soundcard outputs are turned up all the way. Some soundcards let you adjust the output levels of the main outs or individual outputs, so make sure they are turned up and not set too quiet. Again, watch the meters in your soundcard (if it has them) to make sure you're not clipping. Set all soundcard outputs to 0dB if you're not sure. Also, if your soundcard has the option to run the outputs at -10dBv or +4dBu, you might want to try setting it to +4dbu to get a hotter signal coming out of the soundcard.
- Turn up the channel/input gains on the DJ mixer or house mixer you're connected to. Just about every single mixer on the planet has a gain knob that's designed to compensate for signals that are too loud or too quiet. Try turning this up if you can to increase your signal.
- Turn up the PA or sound system to compensate. Ask the sound engineer (if there is one) to boost things on their end, often they still have some gain left on the house mixer they can use to increase the volume of the PA. The good news is that with so many DJs and performers using laptops these days, most clubs are prepared to deal with this issue and the sound person will already have things set up to help you. It pays to discuss this with them ahead of time if you can, ideally at a sound check. For smaller gigs where there is no sound person, check and see if the amps can be turned up more.
- Use an intermediary mixer. Sometimes a small mixer inserted after your soundcard but before the house or DJ mixer can be used to boost the signal more.
- Use plug-ins on your master channel to increase the loudness. I want to stress that this should always be the last resort, as it's by far the one option that alters your sound the most, and rarely for the better. Typically people use things like brickwall limiters to boost the signal until it is loud enough. The downside of this method is that you are reducing the dynamic range of your signal by doing so, and often this was already done to extremes when the song was mastered. So, you might end up with a louder signal, but it comes at the expense of the side-effects of such processes too. Namely this will be more distortion and a dull or weak sound that lacks the punch of the original. Decide carefully if this trade off is worth it to you, just for the sake of a little more volume!

If you do decide to go this route, use the best limiter you have, and adjust it's input gain only as much as needed to get the volume you're looking for. Personal favorites for this are the DMG Limitless and Fabfilter Pro-L limiters. I can't stress enough that you should try all other options before you resort to using plug ins.


As you can see, there's a lot of places in the signal path of a typical DJ or live set where volume can be adjusted. What's more, almost all of them can be changed (for good or bad) by the user, so it really pays to understand proper gain-staging when it comes to your signal. Doing so will insure that you present the crowd with the best sounding audio possible, and in a reproducible manner night after night.

## Digital levels and meters in the studio

Unlike live performers who are often concerned with presenting an audience with the loudest signal possible, in the digital studio our number one goal is to prevent clipping and thus unintended distortion during the production process. More often than not, this means recording our songs without clipping the signals, using the digital medium in the way it was designed for. I've touched on clipping above in the performance section, and the same ideas hold true when we talk about recording in a digital environment. When you exceed the digital medium's upper limit of 0dBFS, the signal will be played back or recorded distorted. Let's look at the places in the production process where this can occur, and see how we can avoid it.

The very first time any audio signal in our studios gets converted to a digital format is when we record something. Whether it's a guitar, a vocalist, a synth, you name it, once we patch it into our soundcards and record it into the DAW, it has become a digital signal. So it makes sense to make sure that when we capture this information as a digital file, we do so the right way. Initially this is pretty easy, we check to make sure that we're not clipping the inputs on the soundcard, using either the soundcard's meters or the input meters of the DAW.

Unlike analog recording where there was some benefit to recording a loud and hot signal, in digital audio we are not aiming to have the peaks of the recording near 0dB. In fact, it's probably safer if we again leave a little bit of a safety net here, and try to keep the peaks of the signal we are recording between -6dBFS and -12dBFS. This ensure that no strays peaks accidentally hit 0dB and are clipped. But aren't there downsides of recording a quieter signal? Actually there aren't, and we have to look at a digital format (in this case a wav file) to understand why.



When you record an audio signal as a wav file, the analog signal is converted into a format that is represented by two attributes; loudness and frequency/pitch. Loudness is represented by a digital file's 'bit-depth', and frequency by it's 'sample-rate'. So when you see a wav file labelled "16bit/44.1kHz", this means there are 16 bits representing how loud or quiet a sound is, and the file is capable of holding frequencies from 0Hz to

22,050Hz (which is half the sample-rate, read on). Let's look at this some more, starting with bit-depth.

The bit-depth of a file directly determines the loudest and quietest sounds it can represent, and this is called the dynamic range. I'll spare you the maths involved, just know that each bit is equal to about 6dB of dynamic range. Therefore an 8 bit file can represent sounds with a 48dB dynamic range. The loudest sound possible is 0dB (this is true no matter what the bit-depth), and the quietest sound we can record is -48dB.

In a 16bit file, we have 96dB of dynamic range, which means we can theoretically record sounds down to -96dB. A 24bit file, as you might guess, has a dynamic range of 144dB. Again, the loudest possible sound is still 0dB, but the noise-floor is now -144dBFS.

So, which do we use? These days I think everyone should be recording and producing using only 24bit files. True they are larger files than the CD standard 16bit files, but the benefits are worth it. And honestly, even a 5 year old 4200RPM hard drive is capable of playing back


dozens of stereo 24bit tracks, so throughput should not be a concern. What are the benefits? For one thing, any artifacts from digital processing are extremely quiet in a 24bit file, well below what we can hear. And of course we also get the benefit of more bits representing our audio signal. Eventually you'll have to convert back to 16bit before you can burn the file onto an audio CD, but this is simple to do with very few downsides compared to the benefits of working in 24bits.

What about sample-rate then? Again, to spare you the math involved (google Shannon-Nyquist) just know that you can record and playback frequencies up to half of the sample-rate. With a 44.1kHz file, this again means 22,050Hz. In a 96kHz file, this means frequencies up to 48kHz can be recorded and played back. Which to use? Unless you have a high quality sample-rate convertor available, I recommend sticking with the CD standard of 44.1kHz. Converting sample-rates is one of the few digital processes that can directly impact audio quality if not done right, and 44.1kHz is still capable of recording frequencies above and beyond the range of human hearing.

There's a few instances where using higher sample-rates is useful however, so if you do have a good sample-rate convertor and don't mind the extra disk space needed, by all means use 88.1kHz or 96kHz. Also, you will often hear people say that if you're going to use a higher sample-rate, then it should be double the rate of your final destination (i.e. 44.1kHz for CD) as that is easier to down sample (just divide by 2) and thus better sounding. This is not true actually, there is no quality loss due to non-integer downsampling, only more CPU consumption during the actual conversion process. The end result will be almost the same, just FYI.

Long story short, if you're not sure what to use, stick with 24bit/44.1kHz wav or aiff files when producing and you'll be fine! This is exactly what I use in my own productions. And this all leads back to why recording at -12dB is not really a problem, even though it might seem too quiet. If you're using 24bit wav files to record to, and each bit equals 6dB of dynamic range, then even with a -12dB signal we still have 22bits representing our audio file  $[(144-12)/6]$ , more than the CD format it will ultimately end up at. In theory you could even record a signal at -48dB and still have all the resolution of an audio CD. Though I don't recommend you take it that far :)

So, now we've chosen a digital recording format, and ensured we have not clipped the audio while converting it to this format. What's the next place we need to worry about clipping? Typically this will be in the track effect plug ins, especially if you're using multiple plug ins on a single track. You want to make sure that the output of one plug in is not so loud that it's going to clip the input of the next plug in. Most plug-ins have output meters to check this, and some DAWs like Live even have handy meters in-between each plug-in so you can verify this:



Now you might think that the next place we could clip the signal would be on the channel meter in the DAW, for instance raising the volume fader on each track until it's making the channel meter turn red and thus indicating clipping. But you'd be wrong. Yes, it does seem like clipping is taking place when the track channel meter is showing red, but pretty much all DAWs use floating-point math at this stage which means that it's almost impossible to clip a channel in your DAW. The maths involved don't belong here (Again! So many maths with digital audio no?), just know that it's impossible for all intents to clip when making a channel fader too loud.

That said, I still think it's a good practice to avoid do this at all costs. It gets you in the habit of setting up good, logical gain-staging when you work, and you never know when that will help out. Perhaps you have to hand your project off to someone else, or you get a chance to work in an analogue studio where clipping a channel really means you're clipping the channel. Always get in the habit of avoiding seeing red on your meters and you'll be safe no matter where or who you work with.

Which brings us to almost the last place we really need to be concerned about clipping a digital signal, the master channel in your DAW. Unlike the individual channels which are impossible to clip, if you see the signal hit 0dBFS on the master channel peak meter, you are definitely clipping the file. I recommend leaving the master fader at 0dB, and adjusting the individual track faders until you no longer see the signal clipping when you play back the entire song. You can see my Mixdown Guide for more info on this.

There is however one final place to be aware clipping can occur, even if it is pretty rare these days. A lot of software offers the ability to normalize an audio file when we export it from our DAWs. Normalizing is a simple process where the entire audio file is scanned to find the highest peak, and then the whole audio file is boosted by the same amount until that peak is at the maximum value of 0dB.

By itself this is no problem, but some older CD players and digital to analog convertors can actually distort in the analog domain when converting a digital file with lots of consecutive samples at 0dB. It's called Inter-sample Modulation Distortion if you want to research the phenomenon some more. To avoid this, most tracks are mastered or normalized so that their highest peak is instead -0.3dBFS. This slight drop prevents the aforementioned distortion, so it's a safe practice to follow.

As with the performance aspect of this guide, you can see that there's multiple places in the production chain where we could accidentally clip our audio files. Sometimes this clipping is not noticeable at first, other times it might be the result of mistakes at numerous points in the process. One thing is certain though, once a file has been recorded with clipping, there's no way to restore it to the way it was before. It's permanently damaged, and we must live with the results or spend the time trying to mitigate them.

More than anything when working with digital audio, it's important to keep this in mind: We don't need to work with signal levels as close to 0dB as possible. In fact, there's a lot of reasons why we should always leave a few dB's of headroom as a safety net. And these few dB's not being used have no impact on the sound of our recordings and performances. Understanding digital audio and how it differs from analog is the best way to ensure that we sound as good as possible, in the studio or on stage!

**Erik M.**

tarekith@icloud.com

-----  
*Feel free to pass this document on as you see fit, though I ask that you do not modify it from it's current form, and give proper credit. If you see any errors, please let me know so I can correct them asap.*