



More Next Blog»



# After School Video

It's not after school, and it's not a video

## Spotify and audio levels

**Mastering your track for streaming**

Or


**Why does my track sound quiet on Spotify?**

Or

**Are the Loudness wars really over?**

Or

**Seriously... what IS going on with Spotify?**



You've recorded your opus. It's mixed, and you've mastered it. You've compared it to other tracks on your hard drive, listened in the car and on headphones and you're delighted. Finally you get it up on Spotify, only to discover...

### Pages

- [Blog p](#)
- [About](#)
- [Watch](#)
- [The Al](#)
- [Spotif](#)
- [Past S](#)

### YouTube

Afte

### YouTube

Afte

### Bandcam

It's quiet. Wimpy rubbish quiet.

What's going on?

Back in the day, you'd get a mastering engineer for your record before it got pressed into vinyl. He (it probably was a he) would make it sound good, and make sure the stylus didn't fly off the record because it was too loud with great lollops of sub bass. Eventually digits came along and we didn't need needles to stop flying off records. But we did still need stuff to sound good, so the mastering engineers stayed even while their role shifted slightly.



As technology evolved, it got possible to make stuff sound louder and louder, still not go above zero (confusingly, zero db is the maximum level a digital file can be, and you work backwards from there) and still sound passable, and thus the Loudness Wars began. Everyone had to compete with everyone else, and gradually the subtleties and dynamic range (the gap between quietest and loudest sounds) eroded. Things were getting over-squashed and quality declining. Eventually the Sound Gods had enough, and a new measurement standard was introduced called... yes, Loudness. This used clever algorithms to work out not technically how high the peak meters got (which was pretty much always zero in the case of pop / rock / EDM etc), but how loud it actually sounded to the human ear. And if your track sounds very loud under a Loudness-managed regime, it will get turned down - conversely if it's very quiet it will get turned up. So everything now sounds roughly the same level, the lion lies with the lamb and everyone lives happily ever after.

The end.

Well... not quite.

### The Dilemma(s)

If you're listening on a CD, or on iTunes from a download, or on Windows Media Player, or on an iPhone or Android from tracks you own, unless you've set it otherwise the chances are that you're listening to the sound exactly as recorded and mastered. Loud is loud, quiet is quiet, and no levels get changed. However, if you're listening on Spotify, Google Play or YouTube, unless you've been fiddling with your preferences and changed them from the default, you're listening to volume compensated tracks - something somewhere has already analysed what you're listening to and turned it up or down to get it in line with everything else. So there are actually two competing standards going on at the same time, depending on what you're listening on and how its set.

And, actually it's worse than that because each of these streaming services uses a

The Al  
by Afte



1. (Se
2. Lo:
3. Th
4. Gl
5. Th
6. Qu
7. Tin
8. Gu
9. Sol
10. Go

The ASV(

different algorithm and different standard (Apple use something called Sound Check, while Spotify uses a different system called ReplayGain). And all the targets and parameters are different for each company. So the idea that everything sounds the same level is very much a theoretical one. And as if all that wasn't head-shrinking enough, as I discovered in many fascinating ways, there's an awful lot that can go wrong in the process, leading to some hair-raising and speaker-cone shattering outcomes.

### So where does it all go wrong?

Let me count the ways...

First of all you have to make a decision. You want your track to sound great of course, and you probably don't want it to sound quieter than everyone else's. But now you have to decide if you are mastering a track to also sound competitive on CDs and downloads, or if its purely to sound competitive on streaming services. What you do next will depend on that decision, because those are two different processes. It could well be argued that the latter is more important, because you're going to be placed directly alongside other artists all the time, and differences will show up starkly.


But can you have your audio cake and eat it? If you're like me you like listening to your favourite stuff on an iPhone through great headphones at maximum level on a run (and, obviously, couldn't care less how stupid you look). If your master is quiet, you just can't go loud enough on a mobile, as the volume control will run out. So chances are you actually want to sound good and healthy on both. Can this be done?

One thing at a time.


You'll need a meter to assess the Loudness of your music. I use one called Dynameter, which is set up with reference presets for Spotify, Apple and so on to show you what you're aiming for, and it provides clear readings of the two most important figures. It takes short term (PSR) and long term (PLR) readings, which will determine how much your music will get turned up or down by the streaming service. In this world by the way, the lower the number, the higher the loudness, so 1 is deafening and 20 is quiet. If its Spotify, at the time of writing you're looking to get a PSR of 8 and a PLR of 10, and Dynameter's creator Ian Shepherd tells us that the PLR is the most important for streaming services (something we'll return to later).

But don't forget - if you also want to sound competitive on downloads, then you'll be wanting those values as close to the target as possible.

What used to happen until very recently is that you could mix a track, then as part of the mastering process put it into something called a Maximizer such as Waves L2, Oxford Inflator, Ozone 7 and so-on. This would magically make your stuff sound loud without going above zero. Better still the really good ones would add a



POWR Tw




real bit of fairy dust - as long as you didn't go nuts, they'd avoid an overcompressed sound somehow, and just make it shine.

One example I was fond of was the Sonnox Oxford Inflator. It sounded fantastic - and here's where I made my first terrible mistake, on After School Video Club's There's Always Someone.

Oh dear - what's the deal with Oxford Inflator?

Unlike, say the Waves L3-16, the Oxford Inflator doesn't get you up to zero and no further. It flies right past zero and into plus numbers - which you'd be forgiven for thinking can't happen. Well it turns out it can and it does - until you save the file. Then it chops off all those plus numbers, producing a radically different result.



There's even a reassuring "clip 0db" button, which sure makes it look like it's stopping everything at zero. The output meter goes up to 0db - and no further. So everything's fine, right? Wrong. "Clip 0db" seems to be a vague aspiration, not a hard and fast rule. Turns out it is sailing merrily past zero, with no way for you to be aware of it. So what happened to me was I'd do my readings in Dynameter, get the thumbs up, save the track and upload it, only to find it turned down as before by Spotify. How could this be? My numbers were great! When I loaded it back in to check, I got totally different readings - because when the track got saved and all those numbers higher than zero vanished, it effectively reduced the dynamics - back into the Red Zone. Which Spotify thus turned down. Oops. Very embarrassing - it all sounded great. Until I pressed save and closed it.

Inflator does all sorts of exciting harmonic things which I like. But you have to put a brick wall limiter after it (so called because you can set a virtual brick wall which it never lets any audio go over), or else it will just chop off all its good work and make your life, as mine, a misery.

So that mystery was resolved, only to crash right into the next one. Despite my Spotify preferences being set to the default "make all tracks the same volume" in preferences, it sounded like I'd ticked the "make all tracks fly out at completely random levels" button instead. Not only was this default button not keeping tracks roughly the same level - it was markedly WORSE than before.

Some tracks blew me, and my speaker cones, to Kansas.

Oh dear. Again. What's the deal with some commercial tracks still sounding INCREDIBLY loud on Windows Spotify?

Edit

@asv

@asvband

Thu Aug 04

Surely

(https:

code-c

album

(https:

(https:

/tweet?i

(https:

/retwee

(https:

/favorite

@asvband

Mon Aug 14

Album

/cults

as 1 le

(https:

(https:

/tweet?i

(https:

/retwee

(https:

/favorite

@asvband

Thu Jul 21

Day 4,

(https:

#false

/falsen

/79RH

(https:

/tweet?i

(https:

/retwee

(https:

/favorite

@asvband

Wed Jul 20

Dream

/electr

https://

/ulPzC

/hasht


(https:

/tweet?i

— " "



I don't go in for conspiracy theories. We landed on the moon, climate change is an actual thing and terrorists brought down the World Trade Center. But I heard stuff coming out of Spotify that defied all known laws of physics, and I began to suspect that some record companies were able to skew the system. Here enters Exhibit A - the excellent and extremely loud Bulletproof by La Roux.



This comes from a really loud album. 80s synths it maybe, but it sounds crisp, clean and deafening, louder than practically any other album I own. Unsurprisingly, if you put it through any loudness meter, it will go off the charts. So Spotify's clever magic will turn it down right?

Well, no.

Not for me, anyway. When I played it on my computer, it barrelled out with all the force of a freight train. Next to the After School Video Club track, it was a bison next to a mouse. "But", I whimpered, "I thought this was meant to make everything sound the same? How can this be? HOW?!!!" Again, in Spotify's Advanced Preferences, "Make all tracks the same volume" was set in its default ON position, which enables the ReplayGain system (you need to do this to be sure you're hearing what 99% of your audience will). And yet for me it still wasn't working, clearly.

After a few weeks of me writing conspiracy theory books about evil record companies, Dynameter creator Ian Shepherd stumbled across the answer as we worked on the problem together. Turns out when I played La Roux on Spotify, it wasn't playing the streaming version at all. Sneakily it was playing my own local copy, unadjusted, at full ear-bleed level. So all the streaming tracks were being turned down more than before, while anything I happened to own stayed just as it was. The difference between the tracks just got even bigger.

#### Blog Arch

▼ 2016 (

▼ Sep

Lov

► Au

► Jul

► Jun

► Ma

► Ap

► Ma

► Fel

► 2015 (

There wa

Hidden in preferences was a setting saying Spotify could import the database from my own Windows Media Player and iTunes libraries. I could be forgiven for missing that these were switched on, because none of the tracks I owned appeared in my Spotify library - just the streaming artists I'd followed. So it was effectively invisible. Its only function, as far as I can work out, was that it would grab any relevant local copy of a file it could get its hands on, and then blow your head off with it.

Thanks for that.

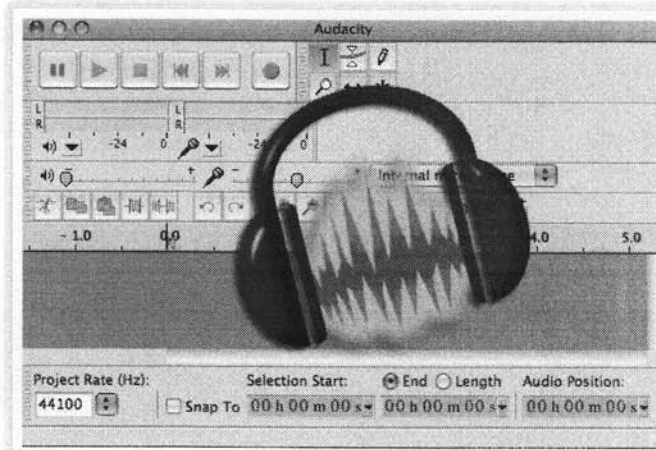
The solution to that little nightmare - deselect "import iTunes / Windows Media Player libraries" in Spotify Preferences.

There, I've just saved you a month of madness. You're welcome.

So now everything's OK?


We're barely getting warmed up.

My next curve-ball was that I found a free plugin to determine any track's ReplayGain level, which worked with Audacity, the free audio editor. This would be very handy - with one click, anyone can see exactly how much a track will be turned up or down by Spotify. Rather than wait 2 weeks for your track to be uploaded only to be disappointed, you can be disappointed instantly instead. Progress!



Or not, as it turned out. While the first few tracks I tried seemed to give sensible results, I'd find the odd anomaly. Sometimes the plugin would return results that seemed quite bonkers. Ah, but could this explain Spotify's occasionally odd results? Should it be trusted?

Furthermore, it returned entirely different results to Dynameter (remember Dynameter and its all important PLR level?) While some mixes correlated well, others were out by quite some margin. Which one to believe? How could you reliably predict what Spotify was going to do to you? I asked both Ian Shepherd at Dynameter and the developer of the ReplayGain plugin. They both said they couldn't be sure what really went on at Spotify.




Let's just take a moment to appreciate that - in 2016, it appears to be impossible to find out what level any track would be played at on the world's biggest streaming service. Nobody really knows - not even the experts.

**This is terrible! Tell me there's a solution!**

There's a solution.

It was time to break out the ol' scientific method. Burning the midnight oil, I fed endless tracks into every measuring device I could find, both before and after they'd been Spotified. I'd watched Mythbusters, and I knew the only difference between science and screwing around was writing stuff down. I wanted science. I wrote stuff down.



Then I looked for any correlation between what these meters said Spotify should do, and what Spotify actually did. The first thing I discovered is I could safely throw out the Audacity ReplayGain plugin - after a few tracks it was clear it was a random number generator. OK, so how about Dynameter's PLR reading? Uh-uh. On its own that turned out to be a poor predictor of what Spotify did as well. That was strange to me because, as I'd understood it, this was supposed to be the most important metric for streaming services. *[Update 16/6/16 Ian Shepherd points out that Dynameter was designed to help you optimise the dynamics of your music in general, rather than make raw predictions - however, prediction is precisely what I was looking for].*

However, my tests suggested that the PSR reading could be more useful for what I wanted - there definitely seemed like a correlation between how much you would get turned down by the Min PSR, rather than the PLR. That said, the PLR was a good gauge to how loud you'd sound BEFORE Spotify did its thing, and that was also important. So in fact, to make a more solid prediction of what would happen, it looked like I could use some combination of the two. After some scratching around, I found... a magic formula.

It seemed to work. I'd run a track across Dynamter, use the formula to predict what Spotify would do to it, then play it out of Spotify to see if it was right. And it was! Time after time it came within 1db of my magic formula's guess.

**Tell me! What's the Magic Formula?!**

Here it is:

# Playback LUFS=Min PSR-PLR-8

Let's go through what these elements mean, and how to use them.

**Playback LUFS** - this is the Loudness figure of the final track. On Spotify, the loudest you can achieve is around -12db to -11db (one track hit -10.5, but that's the absolute maximum I've found). Our track was getting -13db, so was sounding quiet.

**Min PSR** - this is a property of your master before it goes near Spotify, giving the most extreme short term reading during the track. This seems very important in Spotify's adjustments, and the primary way it makes its decisions.

**PLR** - this is the long term measurement. This appears to be a good metric for gauging how loud something really sounds going in. The louder the song measures, the lower the PLR figure - the lowest I've seen is 6. It seems to me that this is the figure that Spotify SHOULD be using, but isn't (I'm told YouTube correlates well with the PLR, by the way).

**8** - this is the specific calibration number I found for Spotify, around which it seems to make its adjustment.

### What about Nugen's Master Check?


This has just been released as this page goes live. In theory it's the dream ticket, just telling you how much your track will be turned up or down on all the streaming platforms. In practice, it didn't produce a good match for me on Spotify - it was suggesting everything would be turned down by about 4db more than it really would. Which is a lot. I've contacted them, hopefully this is something that can be fixed. *[UPDATE 16/6/16 - Nugen agree that the current version isn't working correctly with Spotify, and so hopefully a fix will be forthcoming.]*

### So how can I master a track to get a good final LUFS figure on Spotify?

Well that's the big question, isn't it?

There are a gazillion things that affect the perceived loudness of a track beyond what a bunch of numbers will tell you. But the way Spotify is currently set up steers you in a particular direction. Look at the formula - the thing to keep an eye on is the relationship BETWEEN the short term loudness and long term loudness.

Let's say your track has a PLR of 13 - nicely above Spotify's target of PLR 10. Your overall level going in is quite modest, so you'd expect Spotify to be very happy and let it pass through unscathed, or maybe even crank it up. But things get excitable in the final chorus, and that gives you a Min PSR reading of 5. Oh dear. According to my formula, you start at (PSR) 5, take away (PLR) 13, take away 8. You will likely end up somewhere around LUFS -16 - which is, frankly, pathetic, at under half the perceived loudness of everything else. However, if your song is PLR 11 going in and you keep your Min PSR no





lower than 8 as Dynameter's manual suggests, the equation predicts  $8 - 11 - 8 = -11$  LUFS - bang on Spotify's target value.

So if the gap between PSR and PLR is high - say 5 or more - you'll likely sound relatively quiet (as our first masters did). If that gap is low - 2 or so you'll sound relatively loud. All this is with the necessary caveats for the general mix - La Roux is mixed to sound gobsmackingly loud with a PLR of 6, so even once turned down by 5db, it still sounds pretty punchy.

One other important Caveat - Ian Shepherd agrees that this formula works for loud tracks, but has found it doesn't always work so well at much higher PSR / PLR levels (ie quieter, more dynamic songs). So it's not the absolute literal truth in every case, but a reliable gauge for those who want to sound competitive on Spotify. *[Update 16/6/16 Shepherd confirms my findings, saying that consistently low PSR tends to cause music to be turned down, whereas high PLR can prevent music from being turned up instead. These are important but subtle distinctions and something he agrees need to be spelled out more clearly in Dynameter's documentation]*

### So should I change the way I mix or master?

I've found if there's a technique to be wary of it, it's this - don't use maximizers of any kind as blunt instruments, even if they appear to be doing a good job. If you have a track whose waveform is nicely under 0db peak except for one stretch at the end where everything slams into the maximizer, you'll probably get heavily punished, because there will be a huge gap between short and long term loudness. In theory this can even work in reverse - if you have a quiet stretch that has very restricted dynamics for some reason, that might also skew the end result, but this would be far less common as it's a maximiser's job to achieve exactly this on the loudest passages.

When you look at a regular waveform of your track, the thing perhaps more than anything else you're looking for is no stretch of it that looks like a flat line with everything under it a big block of colour. That's your PSR going into low values, which results in the music being turned down when played online. Ironically, if you really want, you COULD push both into very low numbers - you'll get turned down, but it's louder going in (although Ian Shepherd would point out you're wasting "loudness space", which could potentially give your song more punch and impact). As long as the gap between PSR and PLR isn't too great, you'd be ok. So then it's the case of do you like the effect of maximising, and is it really important to you to sound loud on downloads?

If your master is looking like it's quieter than others once it's been through the equation above, chances are you'll need to go back to the mix and work on that some more, and then try mastering again, using plenty of tools besides maximisers.

Pay close attention to your overall frequency range. The track I learned on, There's Always Someone, had a few issues - too much competing stuff in the lower mids, not enough higher bass frequencies, and a lack of zing at the top end. The mid range generally is important to sounding loud - don't be tempted to get that old graphic eq smiley shape of all bass and treble. Of course this is good mixing and mastering technique anyway - the Loudness measuring standard is just encouraging you to use it.

Other tools to play with include Exciters (with care) and, at the mix level, bus compression and Parallel compression. I think there's some truth in the idea that compressing individual busses enables a more dynamic way to get perceived loudness than a catch-all at the end of the chain, maybe with a bit of glue at the end too. Parallel Compression can be useful if used sparingly and with care, because you can get a whole mix running away from you if you go too nuts.

On There's Always Someone, here were some techniques that made it into the final mix, to produce a result hopefully clearer, less muddy, more energy in the choruses and louder-sounding overall. Not all are applicable to other tracks, but they're examples of what worked well in this case:

Overall careful volume shaping to keep the choruses from peaking considerably more than the rest.

Bus maximising on the beats group (leaving this to final maximising makes it work too hard and you get a lower dynamic range reading), and reducing snare into it

Parallel compression on guitars in the choruses

Compression, EQ and Waves MaxxBass on the main synth bass (MaxxBass adds higher bass frequencies that my mix was overall lacking - less flop, more chest)

Waves MaxxBass on the 808 kick drum

General EQ to add some >10k on guitars, vox and beats

More EQ easing back lower mids on guitars and some vox

More low end on some guitars


And remember - all this is BEFORE it goes to mastering.

Here's a short before and after comparison:

### How Spotify changes the level of your track




Here's how There's Always Someone ended up on Spotify:



**So - can tracks sound loud on both Spotify and iTunes?**

The sad fact is, if you do everything right and you sound terrific on Spotify, you may never get to same levels on your iPhone that you once did. Indeed, La Roux's second album, the tellingly titled *Trouble In Paradise* (not sure the lyrics are all about Mastering, but hey), is mastered much quieter than the guns-blazing debut, something that annoys me on a run - it sounds flat by comparison. Clearly there was a dramatic change of policy. It should be some compensation, however, that if you've had to work that much harder on a mix to make it competitive, that it will sound clearer and better than it did - even if it is a couple of dbs quieter overall than what those Maximisers gathering dust in the corner would have given you.

Hopefully.

 +1 Recommend this on Google

**6 comments:**



Thomas Mavian 14 June 2016 at 20:01

Just brilliant info, thanks!

Reply



Mikko Lohenoja 18 August 2016 at 20:36

Hi!

I'll copy-paste my comment from Production Advice page, as I think you also may be interested in this:

The answer lies in the nature of the ReplayGain algorithm. It measures the 50 ms EQ-weighted RMS blocks, and uses the 95% highest value as the ultimate reference. In ITU/EBU Loudness terms, this practically means it's somewhere at the Momentary Loudness maximum value. In music, Momentary and Short-Term are close together so you can also predict this from PSR values. But for speech, ReplayGain puts the integrated ITU loudness always lower, since speech has larger ratio on Momentary and Short-Term values.

Here are two good examples to measure with Spotify normalisation on. Their absolute, full scale Momentary Loudness values are quite the same, only 0,55 LU apart, but Short-Term even 2 LU apart. Integrated loudness values are very different, almost 7 LU apart.

Leonard Cohen: Famous Blue Raincoat

Max Momentary Loudness -8,34 LUFS

Max Short-Term Loudness -11,28 LUFS

Integrated Loudness -16,98 LUFS

Sex Pistols: Anarchy In The UK

Max Momentary Loudness -8,89 LUFS

Max Short-Term Loudness -9,28 LUFS

Integrated Loudness -10,10 LUFS

If you want to be sure and test out Spotify loudness prior to submitting music, you should do it by using some of the many ReplayGain utilities out there that show the normalisation value for this particular algorithm, or alternatively use an EBU meter and try to keep the Peak-To-Momentary Loudness ratio at about 8.

Hope this helps.

All the best,

Mikko Lohenoja

Helsinki, Finland

Reply



Adam Hutnik 22 June 2017 at 21:57

*This comment has been removed by the author.*

Reply