



Loudness Normalization

Can you turn the volume down?



In this article, we'll take a closer look at loudness normalization, the technology behind it, and its application.

Version 1.0
Thomas Heller

Date 2024-09-18

Loudness normalization provides the luxury of not having to adjust the volume for each audio track that comes on. But you may ask yourself from time to time “Has the last track just ended loudly and the current one is starting softly or does my device not work correctly?” Questions like these will be answered by the tests presented in this post. Here, you get the information and tools to test the loudness normalization of your system and, as a bonus, you learn what target loudness your system is set to.

What is the Difference Between Loudness and Volume?

The loudness of an audio track is calculated over the whole track and saved as metadata in the track next to the audio data. The volume is what the player makes of the audio data and the metadata in combination with the amplifiers of the device that plays the audio track. Loudness normalization means that the loudness metadata is read, and the decoder adjusts all input so that the average loudness stays the same. Broadcasters do this with their programs all the time, but at the sender position. The advantage of modern audio codecs is that they can do normalization at the receiving position. This is much more flexible because it allows the receiving position to set the correct loudness in the receiving device based on the device itself. This means the same audio track is rendered differently on a smartphone with small speakers than on an AVR with high-level speakers.

Who Does Loudness Normalization?

Audio is represented as wave form. This wave form is then played by the loudspeakers that moves air that moves your ear drum. There are different ways to represent audio. Engineers like to talk in terms of the logarithmic unit dBFS (decibels relative to full scale) or LUFS (loudness units relative to full scale) and if two loudness levels are compared, just dB is used. dBFS is used for peaks while LUFS is used for loudness and describes mean energy. To understand this article, it is enough to know that the maximum level for most audio codecs is 0 dBFS, but in practice, the maximum used is -1 dBFS to have some headroom. A smartphone is a low-audio-quality device and should be set to an output range of 15 dB between the average audio and the maximum audio. That is why it has a target loudness of -16 LUFS (which is 15 dB below -1 dBFS). An average TV has enhanced audio (compared to a smartphone) at around -24 LUFS and a high end AVR or soundbar has an out range of -31 LUFS. This means that the loudest part of a symphony is much louder than the average loudness of the symphony if played on an AVR and just a bit louder if played on a smartphone. Just for clarification: the audio does not have to reach the maximum loudness, but it can use it if needed.

Summary

Loudness is metadata that is written into the audio track at the stage of creating the audio bit stream.

Target loudness is a device-dependent setting that the manufacturer of a device sets in the device software.

Volume is what you hear. It's a result of the interactions of loudness, target loudness, and metadata that you can amplify through your volume settings.

Does the Normalization of my Device Work?

“But does this normalization functionality work in my device? And where is my device placed in this range from low quality to high quality?” That’s what our tests are for. There is a PDF online with detailed instructions. In this article we will give an overview of the idea behind the test.

All tests can be done at normal volume by listening or through a microphone recording, depending on how precise you want your results to be – our experience is that listening only shows deviations >2dB.

Preparations

Before starting a test, you should check your setup. Especially if you have a setup that comprises multiple devices connected via HDMI. It's important to know which device creates the audio from the audio track. This creation of audio from data is called decoding and is where the target loudness is used, and the loudness normalization runs. If you have a source device like a Set-Top-Box or TV passing the audio through to an AVR, the audio menu of the source device will say something like "passthrough" or specify a codec. If it says PCM, the source device does the decoding – this is not ideal, since an AVR has fine-tuned settings for its decoding.

Loudness Normalization Check


To check the loudness normalization of an audio codec, play the three files encoded with this audio codec (labeled with -16, -24, -31) and record/ listen to the output volume. If it is at the same level for all files, the normalization works fine. Otherwise, there may be options on your device or in the app on your device to activate loudness normalization. If this is not the case, ask your device's manufacturer about its loudness normalization options.

Target Loudness Check

To check the target loudness of your device, the proposed test uses a trick. AAC-LC is an early audio codec that succeeded mp3 and usually contains no loudness metadata. Therefore the (-16, -24, -31) files will be played at different volume levels. Since there is no loudness metadata, the decoder has no way to know how to shift the waveforms and plays them as they are. To check the target loudness of your device and codec, you can compare AAC-LC with a different codec. For example, if the -31 file of AAC-LC and MPEG-H Audio play out with the same volume, the target device class is at -31. In that case, the -24 AAC-LC will be louder and the -16 AAC-LC even louder. For MPEG-H Audio, the -24 and -16 files should play at the same volume if the loudness normalization works correctly. Professional devices should have the same target loudness for all codecs or vary just a bit. However, this may not be the case: The target loudness of -31 LUFS for MPEG-H Audio, for example, does not mean it is the same for AC-3 or E-AC-3.

Example

To understand both tests, you can play around with the test streams on your PC first. You can use the software "Audacity" and set the recording to "Windows WASAPI". Then you can select a recording option with "(loopback)" in the name, so that the software records what is put out on your system. For example, I select "Loudspeaker (loopback)" on my laptop to record everything that is played on the laptop loudspeakers, independent from which application the sounds originate from. With a right-click on the scale, you can set it to show dB and the mouse wheel helps to zoom in on the relevant area around [-10; -40] dB. Now open a player and play the AAC-LC -16 item and, for example, all AC-3 items followed by the AAC-LC -24 item. Afterwards, you can repeat this using a different player. For me, the results look like this:



What Happens in this Image?

The first player was the Windows media playback. Audio items 2, 3, 4 are at the same level. This is the effect of loudness normalization. Therefore, the conclusion is: The normalization works fine. The target loudness seems to be between -16 and -24, since the surrounding AAC-LC item 1 is higher and item 5 is lower than items 2, 3, 4; therefore, the estimation is that it is at -20.

The second player, in this case the VLC player, fails the normalization test since items 7, 8, 9 have a different loudness and therefore the conclusion is that no target device loudness is set.

It can be concluded that you can see what the pure wave form of items 7, 8, 9 is like. In combination with loudness metadata, they would be on the same level as items 2, 3, 4. And the height of this level is determined by the target loudness of the device, which is at -20 LUFS.

Note: The measured dB values are relative and depend on the volume of the device. Therefore the -20 LUFS is not determined by the dB scale in the picture on the right, which shows 15dB, but by the relationship of the files 2, 3, 4 to files 1 and 5.

Disclaimer

INFORMATION IN THIS DOCUMENT IS PROVIDED 'AS IS' AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

All rights reserved. No part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from the publisher.

Product and corporate names may be trademarks or registered trademarks of other companies. They are used for explanation only, with no intent to infringe.

Copyright (c) 2024 Fraunhofer-Gesellschaft zur Foerderung der angewandten Forschung e.V. as legal entity for its Fraunhofer Institute for Integrated Circuits IIS ('Fraunhofer').