EBU – TECH 3343

# Practical guidelines for Production and Implementation in accordance with EBU R 128

**Supplementary information for EBU R 128**

**Status: Version 1.0**

Geneva
February 2011

[*] Page intentionally left blank. This document is paginated for two sided printing

# Conformance Notation

This document contains both **normative** text and **informative** text.

All text is normative except for that in the Introduction, any § explicitly labelled as 'Informative' or individual paragraphs that start with 'Note:'.

**Normative** text describes indispensable or mandatory elements. It contains the conformance keywords 'shall', 'should' or 'may', defined as follows:

| | |
|---|---|
| 'Shall' and 'shall not': | Indicate requirements to be followed strictly and from which no deviation is permitted in order to conform to the document. |
| 'Should' and 'should not': | Indicate that, among several possibilities, one is recommended as particularly suitable, without mentioning or excluding others.<br><br>OR indicate that a certain course of action is preferred but not necessarily required.<br><br>OR indicate that (in the negative form) a certain possibility or course of action is deprecated but not prohibited. |
| 'May' and 'need not' | Indicate a course of action permissible within the limits of the document. |

**Default** identifies mandatory (in phrases containing "shall") or recommended (in phrases containing "should") presets that can, optionally, be overwritten by user action or supplemented with other options in advanced applications. Mandatory defaults must be supported. The support of recommended defaults is preferred, but not necessarily required.

**Informative** text is potentially helpful to the user, but it is not indispensable and it can be removed, changed or added editorially without affecting the normative text. Informative text does not contain any conformance keywords.

A conformant implementation is one that includes all mandatory provisions ('shall') and, if implemented, all recommended provisions ('should') as described. A conformant implementation need not implement optional provisions ('may') and need not implement them as described.

# Contents

## Acknowledgements

Although this document is the result of much collaborative work within the EBU's PLOUD group, it is the long-suffering chairman of this group, Florian Camerer, who has authored, collated, enriched and distilled the text into its publication form over many, many weeks and months of effort.

## Dedication

This document is dedicated to two great audio engineers; Gerhard Stoll and Gerhard Steinke.

# Practical guidelines for Production and Implementation in accordance with EBU R 128

| *EBU Committee* | *First Issued* | *Revised* | *Re-issued* |
|---|---|---|---|
| Technical Committee | 2011 | | |

**Keywords**: Audio, Loudness, normalisation, production, implementation.

## 1. Introduction

This document describes in practical detail one of the most fundamental changes in the history of audio in broadcasting; the change of the levelling paradigm from peak normalisation to **loudness normalisation**. It cannot be emphasized enough that **loudness metering** and **loudness normalisation** signify a ***true audio levelling revolution***. This change is vital because of the problem which has become a major source of irritation for television and radio audiences around the world; that of the jump in audio levels at the breaks in programmes, between programmes and between channels (*see footnote 2 for a definition of 'programme'*).

The loudness-levelling paradigm affects all stages of an audio broadcast signal, from production to distribution and transmission. Thus, the ultimate goal is to harmonise audio loudness levels to achieve an **equal universal loudness level** for the benefit of the listener.

**Loudness normalisation is a true audio levelling revolution!**

It must be emphasised right away that this does **not** mean that the loudness level shall be all the time constant and uniform *within* a programme, on the contrary! Loudness normalisation shall ensure that the **average** loudness of the **whole** programme is the same for all programmes; within a programme the loudness level can of course vary according to artistic and technical needs. With a new (true) peak level and the (for most cases) lower average loudness level the possible dynamic range (or rather 'Loudness Range'; *see § 2.2*) is actually greater than with current peak normalisation and mixing practices in broadcasting.

The basis of the concept of loudness normalisation is a combination of **EBU Technical Recommendation R 128** '*Loudness normalisation and permitted maximum level of audio signals*' [1] and **Recommendation ITU-R BS.1770** '*Algorithms to measure audio programme loudness and true-peak audio level*' [2].

**EBU R128 and ITU-R BS.1770 are the basis. Four more EBU Technical Documents provide details.**

In addition to **R 128**, the EBU PLOUD group has published four other documents:

- **EBU Tech Doc 3341** '*Loudness Metering: 'EBU Mode' metering to supplement loudness normalisation in accordance with EBU R 128*' [3]
- **EBU Tech Doc 3342** '*Loudness Range: A descriptor to supplement loudness normalisation in accordance with EBU R 128*' [4]
- **EBU Tech Doc 3343** '*Practical Guidelines for Production and Implementation in accordance with EBU R 128*' [this document] and
- **EBU Tech Doc 3344** '*Practical Guidelines for Distribution Systems in accordance with EBU R 128*' [5]

The Technical Documents about '*Loudness Metering*' and about the parameter '*Loudness Range*' also play an important role for the practical implementation of loudness normalisation. They will be introduced as well and referred to in the relevant sections.

The '*Distribution Guidelines*' close the circle, covering all aspects of loudness normalisation for the distribution of audio signals and addressing the critical links between production and the final recipient, the consumer. As this is a very detailed document in itself it will not be introduced here except for the occasional reference.

At the beginning of these '*Practical Guidelines*' the core parts of **EBU R 128** and **ITU-R BS.1770** are introduced, followed by the **general concept** and **philosophy of loudness normalisation**. The document will then look at **loudness strategies** for **production** and **post-production** (metering, mixing, Metadata, etc.), and for **file-based** workflows, that is, ingest, playout and archiving issues (metering, automated normalisation, Metadata etc.).

A separate chapter will look at **Metadata** in more detail. **Alignment** of audio signals and studio **listening levels** are discussed, and **practical advice** is given for the **transition** to loudness-normalised production (implementation and migration). **Genre-specific issues** regarding commercials (advertisements) and trailers as well as music programmes will be addressed in the last chapter.

These Practical Guidelines are meant to be a '**living document**', where, over time, experiences of broadcasters will find its way into the document, providing additional information and guidance for this fundamental change of the way, audio signals are treated and balanced to each other.

Please note that many standards documents are subject to revision from time to time, including this one. You are strongly advised to check for the latest versions.

## 2. EBU R 128, ITU-R BS.1770

**EBU R 128** establishes a predictable and well-defined method to measure the loudness level for news, sports, advertisements, drama, music, promotions, film etc. throughout the broadcast chain and thereby helps professionals to create robust specifications for ingest, production, play-out and distribution to a multitude of platforms. R 128 is based entirely on open standards and aims to harmonise the way we produce and measure audio internationally.

The basis of R 128 is **ITU-R BS.1770**, the result of extensive work by the International Telecommunication Union. The purpose of that standard was to establish an agreed open algorithm for the measurement of loudness and the true peak levels of programmes. It is a robust standard which has the benefit of a simple implementation. In brief, it defines a "**K-weighting**" filter curve (a modified second-order high-pass filter, see Figure 1) which forms the basis for matching an inherently subjective impression with an objective measurement.

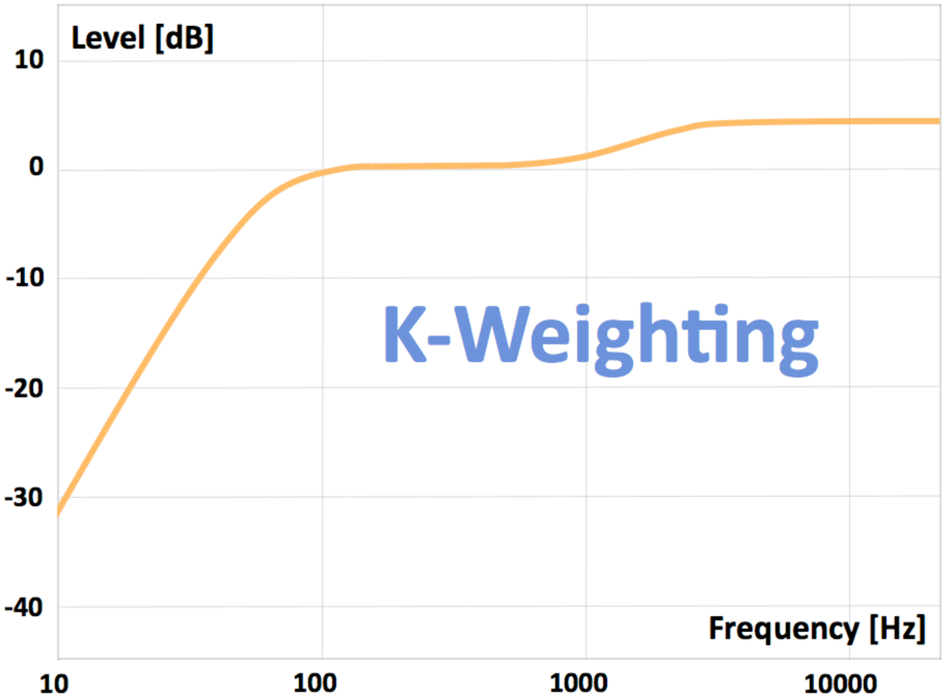


**Figure 1: "K-Weighting" filter curve for loudness measurement**

This weighting curve is applied to all channels (except the Low-Frequency Effects Channel (LFE) which is currently discarded from the measurement; see below), the total mean square level is calculated (with different gain factors for the front and surround channels; see Figure 2) and the result is displayed as "LKFS" (Loudness, K-Weighting, referenced to digital Full Scale), or "LUFS"[1] (Loudness Units, referenced to digital Full Scale). For relative measurements, Loudness Units (LU) are used, where 1 LU is equivalent to 1 dB.

[1] The EBU recommends the use of '**LUFS**' (as specified in EBU Tech Doc 3341). '**LUFS**' is equivalent to 'LKFS' and overcomes an inconsistency between ITU-R BS.1770 and ITU-R BS.1771. '**LUFS**' also complies with the international naming standard ISO 80000-8 [6].

***Low Frequency Effects (LFE) channel***

The Low Frequency Effects channel (the "**.1**"-channel in "**5.1**") of a multichannel audio signal is currently not taken into account for the loudness measurement according to ITU-R BS.1770. This may lead to abuse of the LFE with unnecessary high signal levels. Ongoing investigations try to evaluate the subjective effect the LFE has on the perception of loudness as well as the appropriate way to include it in the objective loudness measurement.

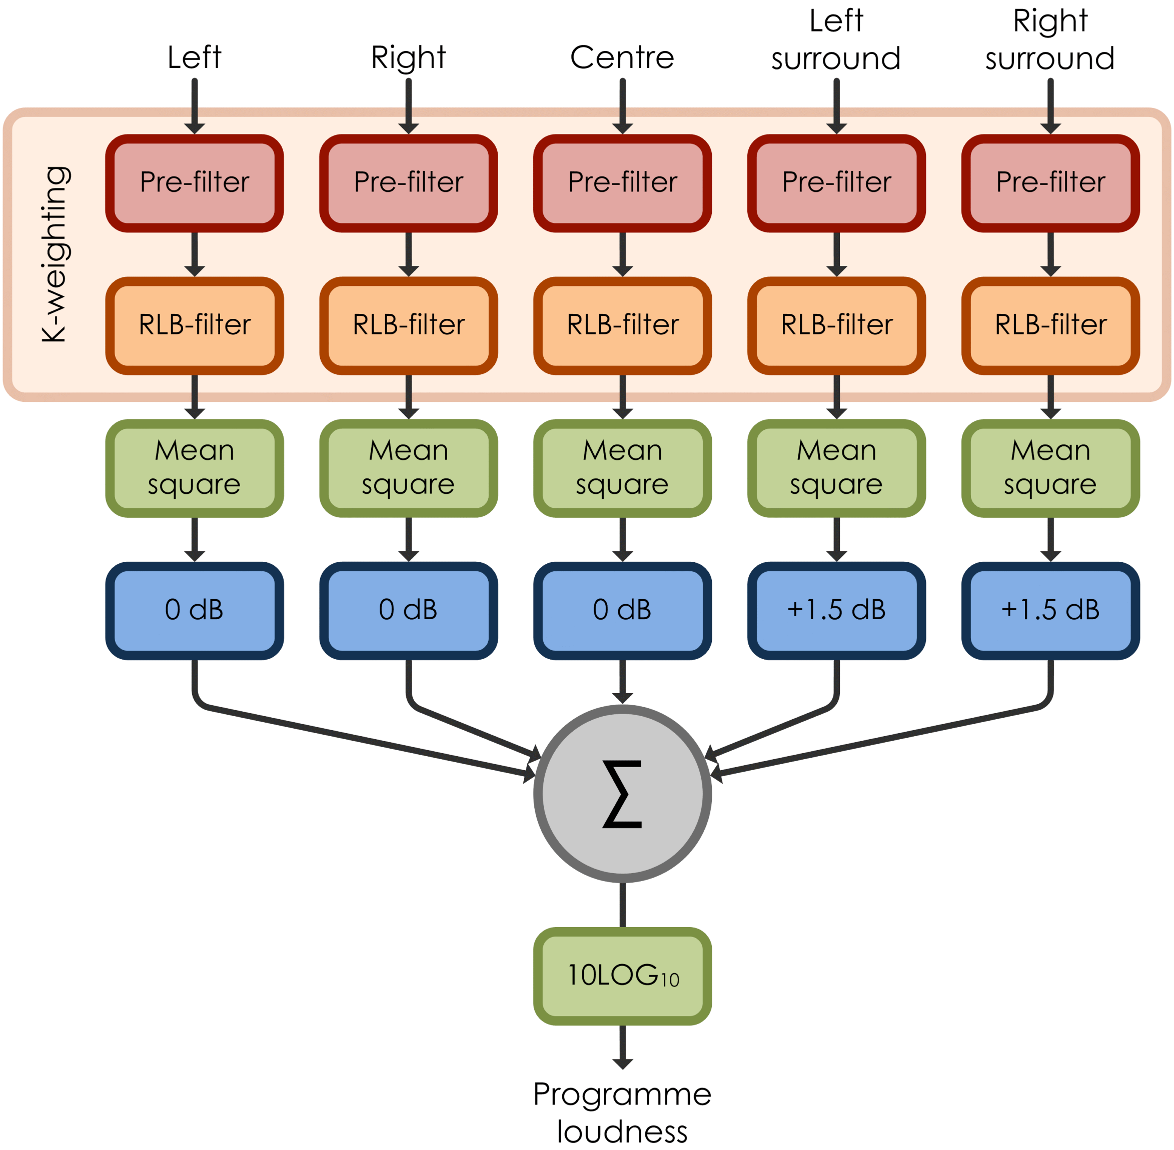


**Figure 2: Channel processing and summation in ITU-R BS.1770**

Whereas BS.1770 defines the measurement method, **R 128** extends it by actually defining a specific '**Target Level**' for loudness normalisation as well as a **gating** method to improve the loudness matching of programmes which contain longer periods of silence or isolated utterances. The EBU's development was needed to accommodate the needs of programme makers, with particular regard to having a means to measure complete mixes (rather than just one component, such as speech or music) and the loudness range of the programme.

To do this, the EBU has specified three new parameters:

- **Programme Loudness**
- **Loudness Range**
- **True Peak Level**

## *2.1 Programme Loudness*

Programme Loudness describes the long-term **integrated** loudness over the duration of a programme[2]. The parameter consists of one number (in LUFS, with one number after the decimal point) which indicates "how loud the programme is on average". This is measured with a meter compliant with ITU-R BS.1770 with the addition of a **gating** function. The gate serves to pause the loudness measurement when the signal drops below a certain threshold. Without this gating function, programmes with longer periods of silence, low-level background sounds or noise will get too low an integrated loudness level value. Such programmes would subsequently be too loud after normalisation.

Following a series of tests, a gate of **-8 LU** (Loudness Units; 1 LU ≡ 1 dB) relative to the ungated LUFS measurement with a block length of 400 ms was agreed to by the EBU (see [7] for more details). This is currently the subject of a revision of ITU-R BS.1770. The tests also confirmed, along with other findings, the choice of the **Target Level** to which every audio signal will be normalised; it is:

**-23.0 LUFS (-8 rel gate)**

**-23 LUFS is the new centre of the audio levelling universe !!!**

***Why -23 LUFS?***

Investigations and measurements of the loudness level of actual broadcasting stations showed an average loudness level of about *-20 LUFS* (with lots of outliers...). Another input to that discussion came from the ITU, in the form of the document ITU-R BS.1864 '*Operational practices for loudness in the international exchange of digital television programmes*' [8]. In BS.1864, a Target Level of *-24 LUFS* is recommended, although without a gating function. Informal tests conducted by members of the EBU PLOUD group have shown that the difference in the loudness measurement with and without the -8 LU relative gate of programmes with a small to medium loudness range are around 0 – 1 LU. **-23 LUFS** with the gate is therefore in many cases almost equivalent to -24 LUFS without a gating method. **-23 LUFS** (with the gate) was therefore considered to be the *lowest possible programme loudness level* without making the transition from an average level of -20 LUFS even more challenging. As -20 LUFS was thought to allow not enough headroom for dynamic

[2] The term '**a programme**' is also used to mean an advertisement, a promotional item etc. For clarity, the advertisements etc. which are placed around and within the running time of what is generally considered to be 'a programme' are treated as programmes in their own right (also individual advertisements within a block are separate programmes); their integration with the longer programmes is thus made easier. Evidently, the makers of either type of programme can have no knowledge of what will be placed with it and so each type has to be considered separately. In this document, the term '**programme**' refers to the programme as completed by Production and not the combination of the programme, interstitials, and advertisements that arrives at the viewer's or listener's receiver within the overall running time of the programme.

content, the decision was taken to settle on **-23.0 LUFS**.

A deviation of **±1.0 LU** is acceptable for programmes where an exact normalisation to the Target Level of **-23.0 LUFS** is not achievable practically (such as live programmes or ones which have an exceedingly short turn-round). In cases where the levels of a programme's individual signals are to a large extent *unpredictable*, where a programme consists of only background elements (for example, the music bed for a weather programme), or where programmes are deliberately mixed lower, the programme loudness level may lie outside the tolerance. Such cases should be increasingly rare, though.

Recently (September 2010), the EBU has submitted the suggestion to the ITU to include the relative gating method into BS.1770. At the subsequent ITU meeting this suggestion was accepted, albeit with a slightly lower threshold level of **-10 LU** below the ungated loudness level. According to the tests of PLOUD, the results with a relative gate of -10 LU are only marginally different from -8 LU. Therefore, once the -10 LU gate is published in the next revision of ITU-R BS.1770, it will also be incorporated into EBU R 128 and the accompanying documents, particularly EBU Tech Doc 3341.

*Details of the gating function need only concern manufacturers. From the user's perspective there will be little difference, but users are advised to keep all their equipment up-to-date to ensure consistency of measurement.*

## *2.2 Loudness Range*

Another major topic was the loudness range which would be needed to accommodate *all* programmes (provided that they don't exceed the tolerable loudness range for domestic listening). The **Loudness Range (LRA)** parameter quantifies (in LU) the variation of the loudness measurement of a programme. It is based on the statistical distribution of loudness within a programme, thereby excluding the extremes. Therefore, for example, a single gunshot is not able to bias the outcome of the LRA computation. EBU Recommendation R 128 does not specify a maximum permitted LRA, as it is dependent on factors such as the tolerance window of the average listener to the station, the distribution of genres of the station etc. R 128 does, however, strongly **encourage the use of LRA** to determine if dynamic treatment of an audio signal is needed and to match the signal with the requirements of a particular transmission channel or platform. More details about LRA may be found in EBU Tech Doc 3342.

**Loudness Range is a generic parameter that helps in deciding if dynamic compression is necessary.**

First experiences at broadcasting stations suggest a maximum LRA value of approximately **20 LU** for highly dynamic material, such as action movies or classical music. The majority of programming will never need to fully use such a high LRA value or, indeed, be able to reach it!

For very short programmes (<30 s) such as commercials, advertisements or trailers, setting a limit for the **maximum** values of the **Momentary** or **Short-term** Loudness Level[3] may provide a better way to control the dynamic properties as a kind of 'second line of defence' (see § 7, § 10).

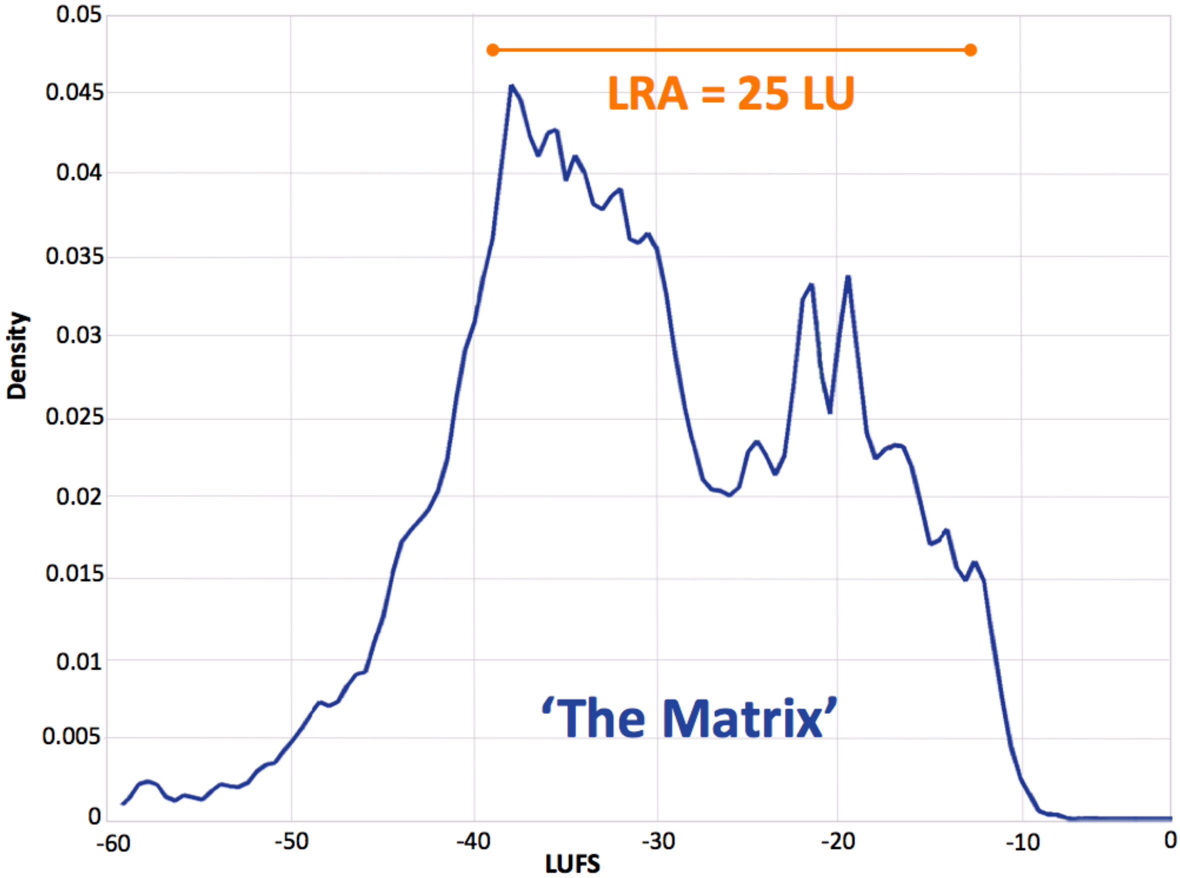


**Figure 3: Loudness Range (LRA) as a result of the statistical distribution of loudness levels**

Figure 3 shows the loudness distribution and LRA of the movie 'The Matrix'; 25 LU is probably challenging for most living rooms...

## *2.3 True Peak Level (TPL), Maximum Permitted TPL*

In Europe, the most widespread metering device was (and still is to a large extent) the **Q**uasi **P**eak **P**rogramme **M**eter (**QPPM**; integration time = 10 ms). With the transition to digital signal processing, **sample peak meters** appeared. While a QPPM cannot display short peaks (<<10 ms) by design, also a sample peak measurement may not reveal the actual peak level represented by a digital signal either.

Digital processing or lossy coding can cause **inter-sample peaks** that exceed the indicated sample level. In broadcasting it is important to have a reliable indication of level across platforms and across sample rates. This meter should indicate clipping, even when the peak lies in between samples, so that the **distortion** that can happen in subsequent Digital-to-Analogue converters, sample rate converters or commonly used codecs can be **predicted** and **avoided**. A sample peak meter cannot do that and is therefore insufficient for use in modern broadcasting (see *Lund, Th.: 'Stop counting samples'* [9]).

[3] **Maximum Momentary Loudness Level** (Max ML) is the highest value (in LUFS) of an audio signal's Momentary Loudness Level (integration time 400 ms). **Maximum Short-term Loudness Level** (Max SL) is the highest value (in LUFS) of an audio signal's Short-term Loudness Level (integration time 3 s).

The **true peak level** indicates the maximum (positive or negative) value of the signal waveform in the continuous time domain; this value may be higher than the largest sample value in the time-sampled domain. With an oversampling true-peak meter compliant with ITU-R BS.1770, those true peaks (unit symbol according to ITU-R BS.1770: **dBTP** - **d**eci**B**el referenced to digital Full Scale measured with a **T**rue **P**eak meter) are now able to be detected. The accuracy depends on the oversampling frequency. It is only necessary to leave a headroom of 1 dB below 0 dBFS to still accommodate the potential under-read of about 0.5 dB (for a 4x oversampling true-peak meter; basic sample rate: 48 kHz).

**Oversampling Peak Meters provide a good estimate for the true peak of an audio signal. Sample Peak Meters don't.**

The **Maximum Permitted True Peak Level** recommended in R 128 will consequently be:

**-1 dBTP**

This is applicable to the production environment for generic linear audio signals. Note that some parts of the chain, such as analogue re-broadcasters and users of commonly used data reduction codecs require a lower True Peak Level. The EBU '*Distribution Guidelines*' (EBU Tech Doc 3344) contain comprehensive coverage of the topic.

### *Summary of EBU R 128*

- The parameters '**Programme Loudness**', '**Loudness Range**' and '**Maximum True Peak Level**' characterise an audio signal;
- The **Programme Loudness Level** shall be normalised to **-23.0 LUFS**;
- The tolerance is generally **±1.0 LU** for programmes where an exact normalisation is not achievable practically;
- The measurement shall be done with a meter compliant with **ITU-R BS.1770** and **EBU Tech Doc 3341** ('EBU mode' - also defining the gating method);
- The parameter **Loudness Range** shall be used to help decide if dynamic compression is needed (dependent on genre, target audience and transmission platform);
- The **Maximum Permitted True Peak** Level in production is **-1 dBTP**;
- Loudness **Metadata** shall be set to indicate **-23 LUFS** (for programmes that have been normalised to that level, as is recommended); loudness Metadata shall always indicate the **correct** value for programme loudness even if for any reason a programme may not be normalised to -23 LUFS.

## *2.4 R 128 Logo*

The EBU has introduced an official logo for R 128, comprised of the numbers **1**, **2** and **8** – forming a happy, smiling face:

The logo may be used (with certain prerequisites) by manufacturers to indicate compliance with 'EBU Mode'.

# 3. General Concept of Loudness Normalisation

## *3.1 Peak vs. Loudness*

The still widespread audio levelling concept of *peak normalisation* with reference to a Permitted Maximum Level (PML; for example, -9 dBFS), has led to uniform peak levels of programmes, but widely varying loudness levels. The actual variation is dependent on the degree of dynamic compression of the signal. In contrast, **loudness normalisation** achieves **equal average loudness of programmes** with the peaks varying depending on the content as well as on the artistic and technical needs (see Figure 4). Provided the loudness range of a programme is within the permitted tolerance, **the listener can enjoy a uniform average loudness level over all programmes, thus not having to use the remote control for frequent volume adjustments any more**.

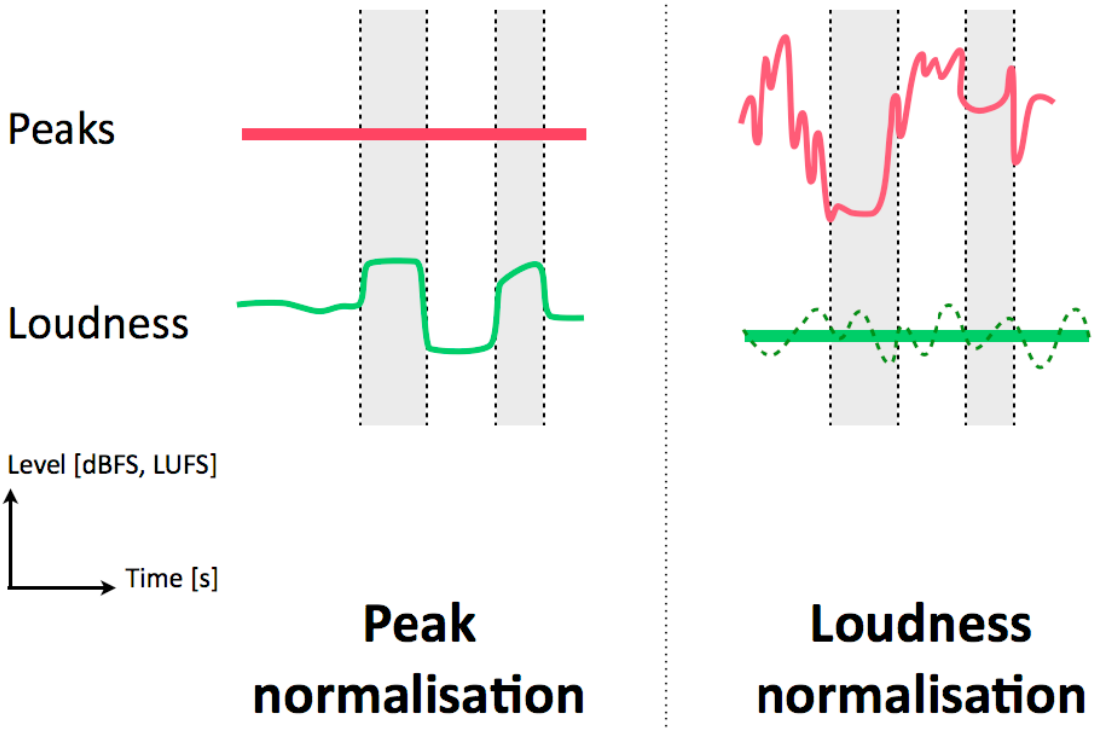


**Figure 4: Peak level normalisation vs. Loudness level normalisation of a series of programmes**

Again, this does **NOT** mean that *within* a programme the loudness level has to be constant, on the contrary! It also does **NOT** mean that *individual components* of a programme (for example, pre-mixes or stem-mixes, a Music & Effects version or an isolated voice-over track) have all to be at the same loudness level! Loudness variation is an artistic tool, and the concept of loudness normalisation according to R 128 actually **encourages more dynamic mixing**! It is the **average, integrated loudness** of the **whole** programme that is normalised.

### 3.2 Normalisation of the Signal vs. Metadata

There are basically *two ways* to achieve loudness normalisation for the consumer: one is the actual **normalisation of the audio signal** itself, so that the programmes are equally loud by design – the other method is with the **use** of **loudness Metadata** that describes how loud a programme is. For the latter, the actual average programme loudness levels don't need to be changed to a normalised value and can still vary a lot from programme to programme. For those with up-to-date equipment, the normalisation can be performed at the consumer's end using the individual loudness Metadata values to gain-range the programmes to the same replay level.

**Equal loudness can be achieved by normalising the audio signal or by using loudness metadata.**

Within the EBU R 128 loudness levelling paradigm the *first* solution is encouraged due to the following advantages:

- simplicity and
- potential quality gain of the audio signal.

The second solution is not forbidden (see also the '*Distribution Guidelines' document, EBU Tech Doc 3344*), but having one single number (**-23 LUFS**) has great strength in spreading the loudness-levelling concept, as it is easy to understand and act upon. And the active normalisation of the source in a way 'punishes' overcompressed signals and thus automatically encourages production people to think about other, more dynamic and creative ways to make an impact with their programme. In other words, the actual *technical* change of the audio signal level through **active normalisation to -23 LUFS** has direct consequences on the *artistic* process – and in a positive way! The production side is thus relieved from fighting the 'loudness war' – an unfortunate and widespread result of the peak-normalisation paradigm.

Nevertheless, it must be stated that both methods can complement each other, they are not to be seen as opponents or a black-and-white view of the same issue. Both approaches are a part of R 128 – but because of the advantages listed above, the **normalisation of the audio signal** is recommended.

**Loudness normalisation of the audio signal is recommended in production because of simplicity and potential quality gain.**

Working towards a common loudness level signifies a **whole new concept** of mixing, of levelling, of generally working with audio. Whereas a peak limiter set to the Permitted Maximum Level (usually -9 dBFS, measured with a QPPM) provided a sort of '*safety ceiling*' where, no matter how hard you hit it, it always ensured the 'correct' maximum level, the loudness levelling paradigm more resembles '*floating in space, with the open sky above*' (see Figure 5).

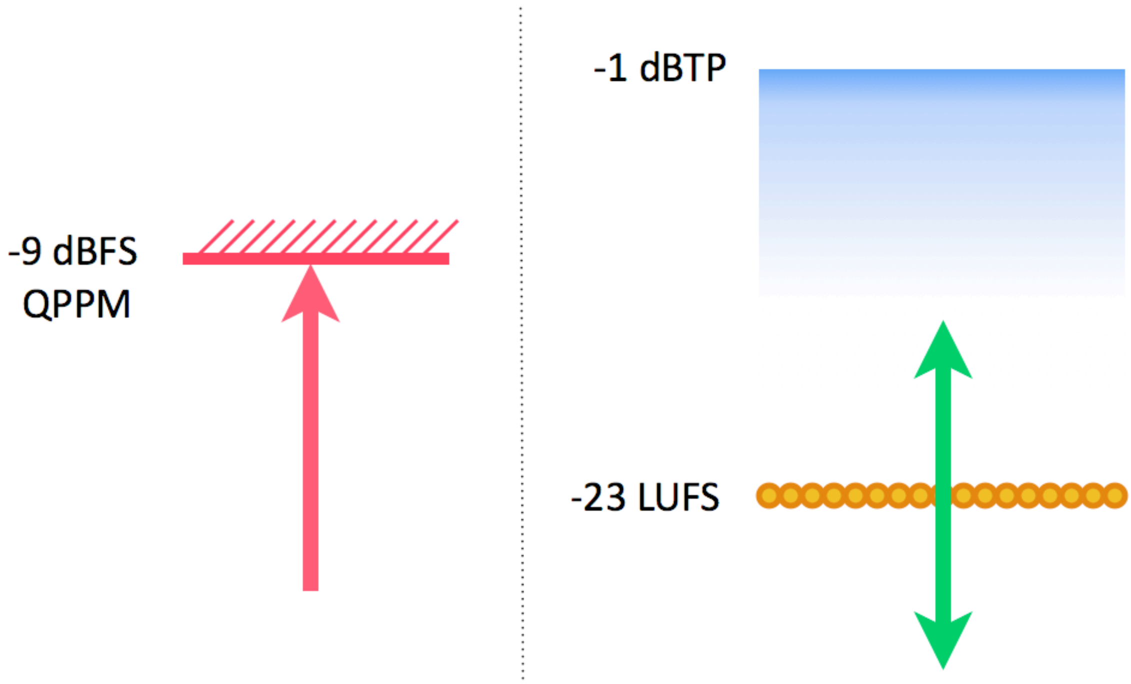


**Figure 5: Quasi-Peak Level normalisation ('safety ceiling') vs. Loudness Level normalisation**

With loudness normalisation and metering, the safety ceiling is gone. This might be intimidating for some, as it was in a way 'comfortable' that one didn't have to watch levels so closely – the limiter at the end of the chain ensured that your output was always tamed. But the side effect was that loudness levels went up, because more sophisticated processors exhibited fewer dynamic compression and limiting artefacts.

**Loudness levelling**, on the other hand, encourages the use of by far the best metering device: **the ear**. This implies more alert mixing and fosters audio quality. Experience of several EBU members has shown that working with the loudness paradigm is **liberating** and satisfactory. The fight for 'Who is the loudest' is gone, overall levels go down, and this in combination with a higher Maximum Permitted True Peak Level (**-1 dBTP**) results in potentially **more dynamic mixes** with **greater consistency of loudness** within a programme. Dynamic compression is again an artistic tool and not a loudness weapon, the **audio quality increases**!

Putting '**mixing by ear**' back on track is a welcome relief and long overdue. The mixer is now encouraged to mix by ear alone (another effect of loudness metering) – after setting levels at a fixed monitor volume (see § 8).

**Loudness levelling encourages mixing 'only by ear' - after setting levels and a fixed monitor volume.**

Downstream of production, the broadcaster is confronted with the need to normalise diverse content originating from different places. Especially during the transition period there will still be many programmes that are not yet loudness normalised. Strategies for these programmes have to be developed, like **automated normalisation** directly after ingest to a playout server or the installation of a safety loudness regulation device at the output of master control to be able to handle, for example, live feeds that are not produced to the target level of **-23 LUFS**.

Issues like these will be covered with additional details later on (§ 4.3, § 6).

### 3.3 Using the Parameter Loudness Range

For the first time it is now possible to *quantify* the dynamics of a programme. In the past, it had to be 'educated guesswork' of experienced audio personnel to decide if a programme would fit into the loudness tolerance window of the intended audience. Using **Loudness Range (LRA)**, the guesswork is over – at the end of the measurement period (usually the whole programme), a single number enables the mixer/operator to decide if further dynamic treatment is necessary.

It is important to understand that it is impossible to define one maximum value of LRA for all broadcasters and all programmes. The **individual maximum LRA** is dependent on the genre(s) (theme channels with very uniform content like news will certainly not have as high a maximum LRA as public broadcasters with a great variety of genres for programmes like, for example, a classical music concert). The maximum LRA value may also be different for distribution platforms like mobile broadcasting as well as different replay environments (*see Figure 6; the distance between the yellow lines indicates examples of different Loudness Range values*). The average listening environment, age of the target audience, 'listening comfort zone' of the consumer and other parameters all influence the choice of a station's maximum LRA values for specific programming. The **Loudness Range Control Paradigm** starts from a generic maximum accepted value of Loudness Range according to the principles described above and adapts this value downstream to comply with technical necessities of individual distribution platforms and replay environments.

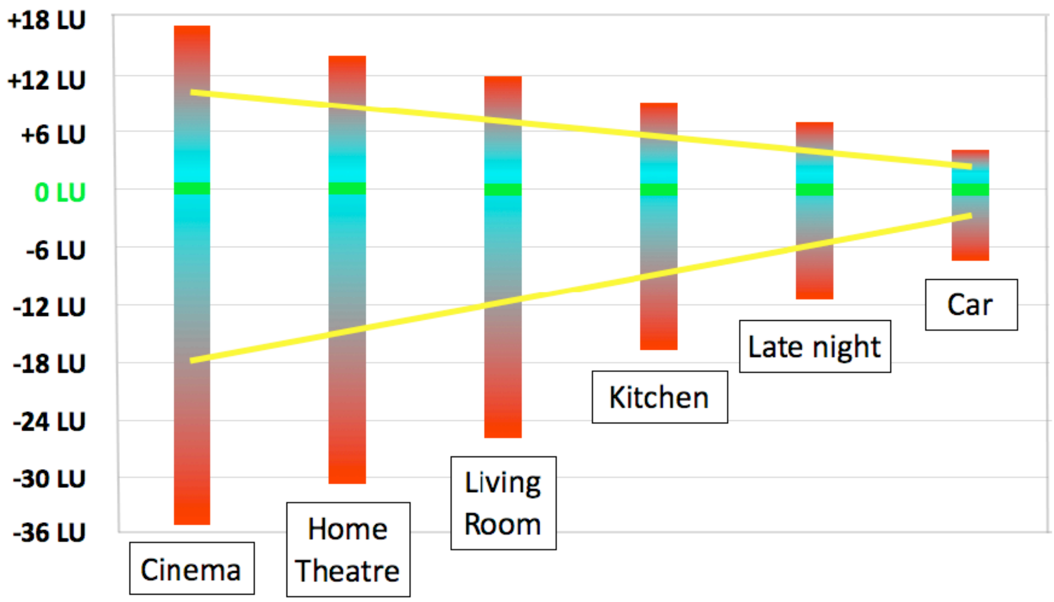


**Figure 6: Different examples for Loudness Range depending on the replay environment**

It is the responsibility of the mixer at the production stage of the programme to determine the programme's LRA; the mixer is advised to respect the paradigm. In environments where human control and intervention are not possible or in place, the measurement of LRA may lead to the use of suitable presets of a dynamic processor that has been configured for the individual genre of the audio signal. Nevertheless, it is advisable to strive for a situation where an audio engineer can influence LRA according to the specifications of the broadcaster, as this potentially increases audio quality.

As a result of the need for different values of Loudness Range, **EBU R 128** does *not* include a maximum permitted LRA value, but instead strongly encourages the **use** of the **Loudness Range** parameter to evaluate the potential need for dynamic range processing according to the different criteria mentioned above.

Loudness Range is also a useful **indicator** of potential *dynamics reduction processes* in a signal chain, performed on purpose or accidentally. If the LRA value of a programme after it has passed through a processor chain is, for example, lower than it was originally, such a reduction process has occurred.

## 3.4 Climbing the True Peak

The third parameter recommended by R 128 concerns the **maximum true peak level** of an audio signal. Having abandoned the peak normalisation paradigm, it is of course **still vital to measure and control the peaks of a programme**, and especially its maximum peak to avoid overload and distortion.

A loudness meter compliant with '**EBU mode**' (see *EBU Tech Doc 3341*) also features the measurement and display of the true peak levels of a programme. Safety limiters to avoid overmodulation will have to be able to work in *true-peak mode* and need to be adjusted to the appropriate Maximum Permitted True Peak Level, in production as well as at the output of master control, at the distribution headend and the transmitter site. Next to the Maximum True Peak Level for generic PCM signals in production (-1 dBTP), further suggestions for different applications are given in EBU Tech Doc 3344 ('*Distribution Guidelines*').

# 4. Strategies for Loudness Normalisation

## *4.1 Production, Post-Production*

Approaching loudness normalisation in these areas offers two possibilities: the first is to keep current levelling practices and perform a level shift afterwards, and the second is to change the levelling habit to **loudness control** and **normalisation** with no or only a small shift needed (Figure 7).

The first approach is more relevant for the early stages of the transition, and it is perhaps especially useful to those who work on **live programmes**. The existing meters, limiters and mixing practices are retained and a level shift is done at the output of the console (after the main meters) to achieve the loudness **Target Level** of **-23 LUFS**.

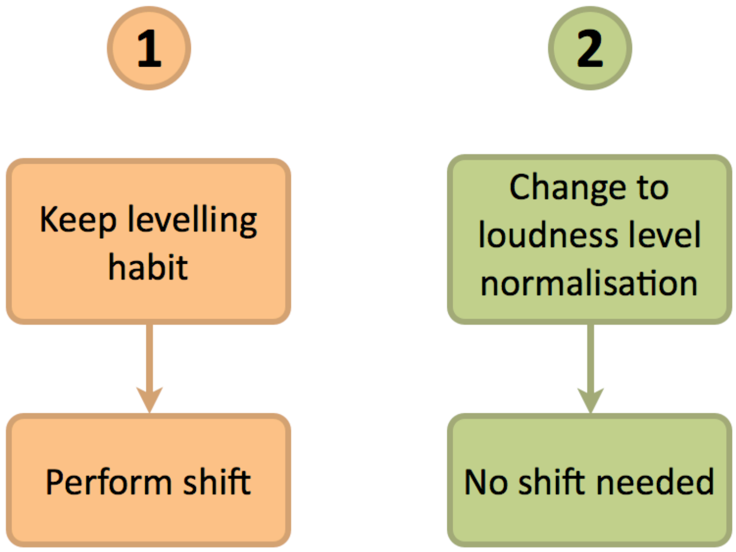


**Figure 7: Two principal working methods to achieve uniform loudness in production and post-production**

A loudness meter is placed after the level shift to enable the engineers to understand the exact amount of shift (which initially is still a bit of guesswork). Using a loudness **meter in parallel** with a conventional meter is in any case a good idea. In this way experience can be gained before actually diving into the loudness-levelling world. Furthermore, using a loudness meter to measure past programmes of the same genre gives good guidance as to where the levels sit.

For programmes that are finished in **post-production** the necessary level shift for approach **1** is easy to perform. Measuring the whole programme in one go, the necessary gain offset can be determined exactly, and in today's file-based world a gain calculation is a very quick and easy operation.

Of course, for **live programmes** it is challenging (if not a matter of luck) to exactly achieve Target Level. Therefore, a deviation of **±1.0 LU** is acceptable for programmes where an exact normalisation to the Target Level of -23 LUFS is not achievable practically (in addition to live programmes, for example, ones which have an exceedingly short turn-round). Early experience at NDR (North-German Broadcasting), ORF (Austrian Broadcasting), and RTBF (Belgian Broadcasting - French part) has shown that it is certainly possible for live mixes to fall within the ±1 LU window permitted by EBU R 128.

**A tolerance of ± 1 LU around target level (-23 LUFS) is accepted. Exceptions in special cases are anticipated.**

In cases where the levels of a programme's individual signals are to a large extent *unpredictable*, where a programme deliberately consists of only background elements (for example, the music bed for a weather programme) or where the dramaturgical intention of a programme makes a loudness level particularly lower than target level desirable, this tolerance may be too tight. It is therefore anticipated for such cases that the integrated loudness level may lie outside the tolerance specified in R 128.

For **levelling solution 1** (keeping current levelling practices) it is likely that in almost all cases the necessary gain shift will be *negative* (attenuation). Therefore an additional step of reducing the dynamic range and/or limiting the Maximum True Peak Level is usually not necessary. The potential attenuation in the vast majority of cases is also the reason why the Metadata solution described in § 3.2 is not advised for solution 1.

**Levelling solution 2** (changing to loudness normalisation right away) is the one that is **recommended** in these Practical Guidelines. Again, after an initial measurement and testing period of past programmes and the installation of a loudness meter in parallel to the meter currently used (usually a QPPM), the advantages of the loudness-levelling paradigm speak for themselves. The greater dynamic range possible will be a welcome bonus for crowd noise, for example, of sports programmes, enhancing the impact of a game for the viewers and listeners. Studio voice-overs that are often dynamically compressed due to artistic reasons (and where therefore the loudness-to-peak ratio will be lower) will be better balanced with more dynamic original location recordings etc. etc.

In what follows, the impact of working with a loudness meter in production and post-production will be examined.

## *4.2 Loudness Metering for Production and Post-Production*

An ‘**EBU Mode**’ loudness meter as defined in EBU Tech Doc 3341 offers 3 distinct time scales:

- **Momentary** Loudness (abbreviated “**M**”) – time window: **400 ms**
- **Short-term** Loudness (abbreviated “**S**”) – time window: **3 s**
- **Integrated** Loudness (abbreviated “**I**”) – from ‘**start**’ to ‘**stop**’

The **M** and **S** time windows[4] are intended to be used for the immediate levelling and mixing of audio signals. Initial level setting may be performed best with the Momentary Loudness Meter, adjusting the level of key or anchor elements (such as voice, music or sound effects) to be around the Target Level of **-23 LUFS**. Of course a mixer has to know at any time how loud the actual signal is, and that is the main purpose of the Momentary and Short-term measurement.

Due to an inconsistency between ITU-R BS.1770 and ITU-R BS.1771, EBU Tech Doc 3341 suggests a different naming convention, complying with ISO 80000-8:

- The symbol for ‘Loudness Level, K-weighted’ should be ‘$L_K$’.
- The unit symbol ‘**LUFS**’ indicates the value of $L_K$ with reference to digital full scale.
- The unit symbol ‘**LU**’ indicates $L_K$ without a direct absolute reference and thus also describes loudness level differences.

Any graphical or user-interface details of a loudness meter complying with ‘**EBU Mode**’ have deliberately not been specified; nevertheless, two scales have been defined: **“EBU +9 Scale”** which ought to be suitable for most programmes and “**EBU +18 Scale**” which may be needed for programmes with a wide LRA. Both scales can either display the relative Loudness Level in **LU**, or the absolute one in LUFS. **‘0 LU’ in ‘EBU mode’ equals the target level of -23 LUFS**. The meter manufacturers in the PLOUD Group have agreed to implement the ‘EBU Mode’ set of parameters, to make sure their meters’ readings will be aligned.

**In an ‘EBU Mode’ loudness meter, 0 LU equals -23 LUFS.**

Many more manufacturers have adopted ‘EBU Mode’ too, or are in the process of doing so. Figure 8 shows a schematic representation of a bar-graph meter with the two EBU-mode scales; Figure 9 shows how a software meter based on a “needle” could look.

[4] ‘**M**’ and ‘**S**’ are commonly used in stereophony for ‘Mid’ and ‘Side’. To distinguish the integration times ‘Momentary’ and ‘Short-term’, the versions ‘**M**$L_K$’ and ‘**S**$L_K$’ (as well as ‘**I**$L_K$’) may be used. ‘$L_K$’ stands for ‘Loudness Level, K-weighted’, and complies with the international naming standard ISO 80000-8.

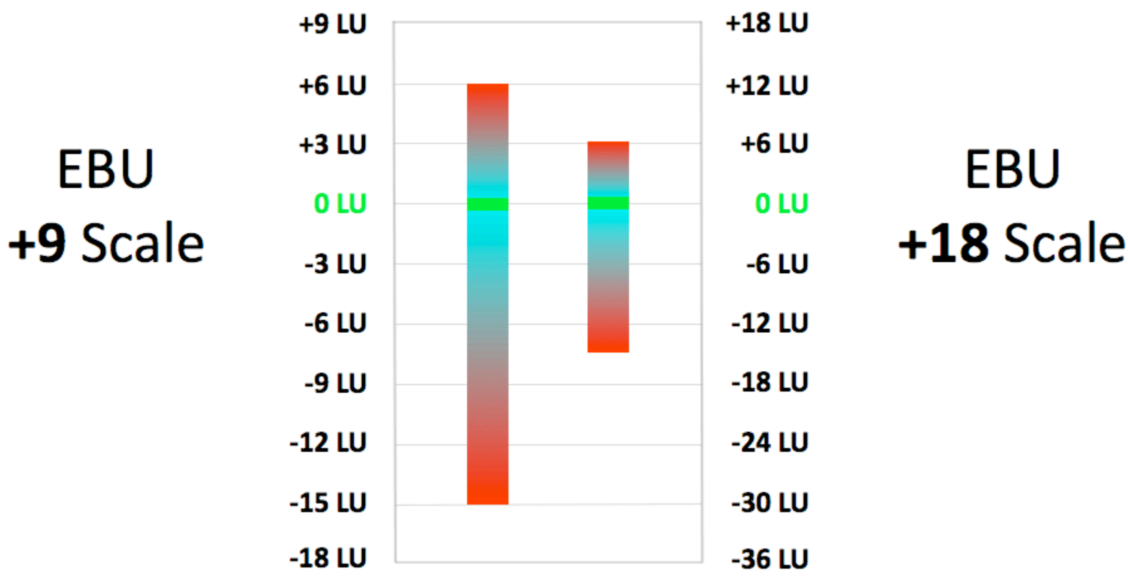


**Figure 8: A schematic representation of the two loudness scales (here in LU) as described in EBU Tech Doc 3341**

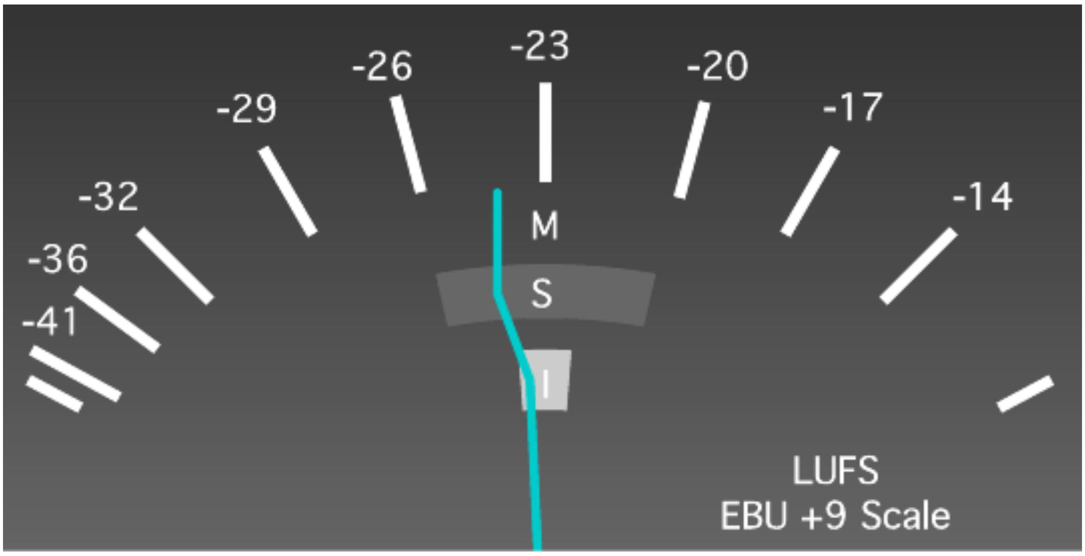


**Figure 9: A schematic representation of an emulated loudness meter with a "bendy needle"**

## *4.3 "Ready, Set (Levels), GO!"*

It is advisable to **set levels** with a bit of **caution** initially, as it is psychologically easier to gradually increase your integrated loudness level during a mix than to decrease it. The gating function also has an influence: bigger level changes are needed to bring the average level down than to raise it. Usually, a slight increase in the course of a programme is also dramaturgically more natural – and an initially "defensive" strategy leaves the engineer room to manoeuvre in case of unexpected or unpredictable signals and events.

Once levels of individual signals are set, and a fixed monitor gain has been established (see § 8), the audio engineer can switch to **mixing only by ear**. Watching the Momentary or Short-term loudness level and an occasional glance at the value of the Integrated loudness level should give enough confirmation that the mix is on the right track towards Target Level. With a numerical readout of the 'I'-value with one decimal point precision or a graphical display of similar resolution, **trends** can be anticipated and the appropriate measures taken. This should be performed in a smooth manner, as too drastic changes will, in most cases, be artistically unsatisfactory.

With the **Maximum Permitted True Peak Level** being **-1 dBTP**, the phenomenon of 'hitting the wall' (meaning the safety limiter operating at -9 dBFS) is now less likely to occur. Used reasonably and with a clear intention, this '*opening of the lid*' together with loudness normalisation to -23 LUFS results in **more dynamic mixes**, in less dynamic compression artefacts like pumping and thus in an **overall increase of audio quality**! Programme makers who favoured dynamic mixes in the past are now relieved from potential compromises because their programme would sound softer than more compressed ones. With loudness normalisation, this compromise is gone. At last!

*The elements of a mix that are most important for a uniform subjective loudness impression are so-called '**foreground**' sounds – like voice, music or key sound effects. Individual sound elements do have a widely varying difference between their loudness level and their peak level. For example, the 'clink' of two glasses when toasting has a high peak level, but quite low loudness level. On the other hand, a dynamically compressed hard rock guitar riff has a loudness level that is almost the same as its peak level! If those two signals are aligned according to their peaks, the guitar riff will be much louder than the clink of the glasses. This example is meant to illustrate the concept, it does NOT mean that those two signals are necessarily to be mixed with equal loudness! The level of individual elements and components (like pre-mixes or stem-mixes, a music-only mix or a voice-over track) in the mix is an artistic decision, naturally, but **loudness metering** can help the mixer with useful visual feedback that actually shows what he or she hears!*

Coming back to metering, at the **end** of a **programme** there are **two scenarios**:

- having exactly hit target level (-23.0 LUFS) or
- having missed target level in either direction

For **live productions**, understandably the second scenario will be more likely. If the actual loudness level is within the accepted tolerance of **±1.0 LU**, then no further action is needed. If the level lies outside this tolerance due to the particularly unpredictable nature of the programme or the rare occurrence of foreground elements, this is still acceptable from a generic production standpoint (as mentioned earlier). Measures may be taken further downstream to 'tame' these cases in the form of **loudness processors** that gradually adjust the integrated loudness level of such programmes in an unobtrusive manner and can act as a sort of '*loudness safety net*'. This must be achieved with a reasonably slow reaction time, so that the inner dynamics of the production are not harmed. *Differentiation* between live and file-based programmes should be possible as far as the individual preset of such a loudness processor is concerned or where in the signal chain the processor is installed. The processor may only be needed for live programmes if the workflow for file-based programmes is already fully compliant with EBU R 128. If a downstream dynamics and loudness processor is situated at the output of the **Master Control Room**, it should be able to be **bypassed** for programmes compliant with R 128. This bypass situation is expected to become the normal way of working the more programmes are 'on target', as the ultimate recommended goal is to tackle and **normalise the audio signal itself**.

In the **post-production** area one is more likely to hit target level because of the very nature of the workflow with more opportunities to redo and change the mix and thus loudness levels. Furthermore, usually there is enough time to perform a complete integrated measurement of the whole programme once it is finished, as well as a subsequent **gain correction**. In a file-based production environment this correction can be performed much faster than realtime. Situations may be common where mixes in the post-production area are performed '*as live*', that is, for example, directly onto a tape with few if any mistakes from voice talents (in case of a voice-over mix). Also, for example, a 1:1 **tape-copy** process with loudness adjustment 'on-the-fly' falls into this category. These situations are then similar to live productions and should be approached accordingly.

Especially in the **transition phase** moving towards loudness normalisation such aforementioned loudness processors downstream will certainly be helpful for broadcasters to adapt to the loudness levelling system and catching possible outliers. It should be the goal of the broadcaster (and also the mixing engineers) to have these processors work as little as possible, as the integrated loudness level of programmes is increasingly within the accepted loudness tolerance. The exact transition scenario, time scale and implementation plan are of course individually different for each broadcaster (see § 9). While this is anticipated, in the interest of the consumer the switch to loudness normalisation should be performed in due time as the benefit for the listener is so substantial.

### *4.4 Loudness Range for Production and Post-Production*

Working with loudness normalisation right away also implies controlling **Loudness Range** (LRA) as the dynamic possibilities are expanded. This is important to ensure an appropriate signal for the intended audience and distribution chain. Whereas in production and post-production a '*generic*' mix may be created (with a relatively high LRA value and a Maximum Permitted True Peak Level of -1 dBTP), different platforms may need a lower LRA value and a lower Maximum Permitted True Peak Level (while keeping the Programme Loudness Level at -23 LUFS). **The system within R 128 appreciates this generic approach with further processing downstream to tailor the signal to individual environments and platforms**.

With the parameter Loudness Range it is now systematically possible to determine appropriate measures for *potential dynamic compression* of a programme to fit it to the tolerance window of the audience or distribution platform. In practice, **overall low-level compression** may lead to satisfactory results (*see Figure 10 as an example*): a low threshold (< -40 dBFS) and a moderate compression ratio (1:1.2 – 1:1.5), together with a long release-time (>0.5 s), ensure uniform compression of the whole signal range.

**Compressing LRA** (example):

- **Low threshold (< -40 dBFS)**
- **Low Ratio (1:1.2 - 1:1.5)**
- **Adjust make-up gain accordingly**

Dependent on the original loudness level, a shift to the Target Level of **-23 LUFS** may be performed in parallel through adjusting the make-up gain of the compressor accordingly.

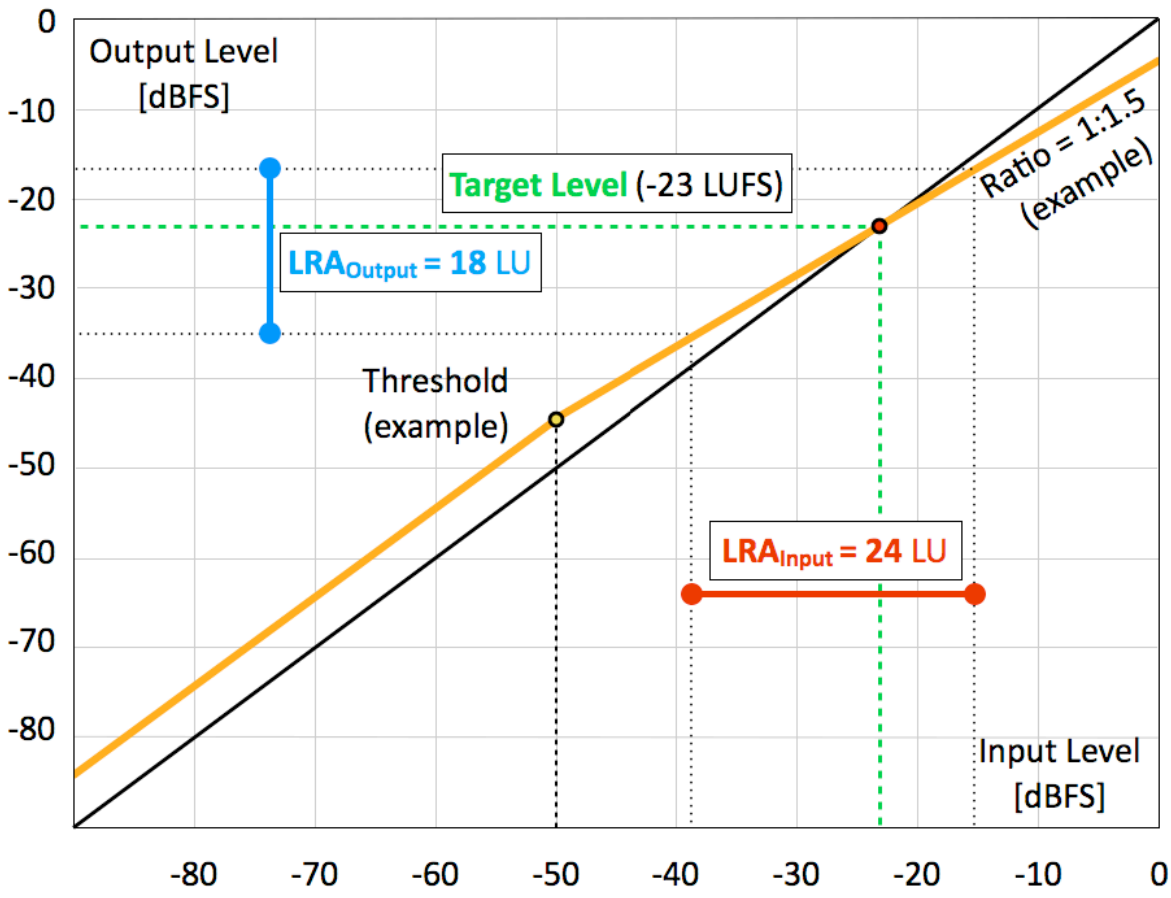


**Figure 10: Example for processing of Loudness Range (LRA) with a compressor with a low threshold (-50 dBFS) and a moderate compression ratio (1:1.5)**

# 5. What to Measure in Production and Post-Production

## *5.1 Signal-Independent vs. Anchor-Based Normalisation*

**EBU R 128** recommends measuring the **whole programme**, independent of individual signal types like voice, music or sound effects (see figure 11). This is considered to be the most generally applicable practice for the vast majority of programmes:

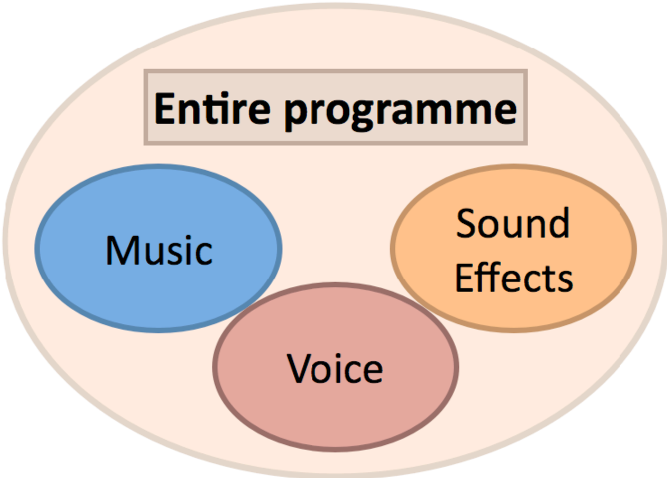


**Figure 11: Elements of a programme**

For programmes with an increasingly wide loudness range (>12 LU, approximately) one may optionally use a so-called **anchor signal** for loudness normalisation, thus performing an *individual gating method*, so to speak. This should be a signal which the producer or engineer wants to be representative for the average loudness of the programme, like speech or a singing voice, a certain part of a music programme in *mezzoforte*, a consistently applied and dramaturgically important sound effect sequence etc.

It must be emphasized, though, that choosing an anchor signal is an **active process** requiring input from an experienced operator. This approach should only be considered after operators and sound engineers have become very comfortable with the concept of loudness normalisation. Performed well, it may help to fine-tune the loudness of wide loudness range programmes according to the chosen anchor signal.

There also exists an automatic measurement of one specific anchor signal in the form of '*Dialogue Intelligence*', a proprietary algorithm of Dolby Laboratories, anticipating that speech is a common and important signal in broadcasting. The algorithm detects if speech is present in a programme and, when activated, only measures the loudness during the speech intervals. For programmes with a narrow loudness range the difference between a measurement restricted to speech and one performed on the whole programme is small, usually <1 LU. For programmes with a wide loudness range, such as action movies, this difference gets potentially bigger, sometimes exceeding 4 LU. Automatic detection of an anchor signal is intended to help identify what should be at Target Level. Like any algorithm for detecting specific signals out of a complete and complex mix, speech discrimination can be tricked – either by signals closely resembling the spectral pattern of speech (for example, certain woodwind instruments or a solo violin) or by speech signals that are too far off the discrimination threshold (for example, certain language dialects). For programmes where these anchor signals are consistently moving around the discrimination threshold, the loudness measurement can also vary significantly if the measurement is performed repeatedly.

For **short** programmes like commercials, advertisements, trailers and promotional items, (automatic) speech normalisation is likely to give unsatisfactory results in the light of the future increase of loudness range and potentially enhanced dramaturgical concepts. In such cases, most international recommendations (also this one) agree on measuring '**all**' by all means.

In any case, broadcasters have to be aware that especially in a file-based environment, where for most content the whole programme independent of signal type (speech, music, sound effects) will be measured automatically, a different strategy might have to be established to treat programmes based on anchor normalisation.

***To summarize:***

It is because of these uncertainties and the fact that speech represents only one part of the whole programme (albeit a very important and common one) that R 128 recommends measuring '**all**' – that is the whole programme, independent on the signal type (such as voice, music or sound effects).

This is supported by the following observations:

- The difference between measuring 'all' and measuring an anchor signal (such as voice, music or sound effects) is small for programmes with a narrow Loudness Range;
- The difference between 'all' and 'anchor' measurements depends strongly on the content of the programme, but can be expected to be bigger if the Loudness Range is bigger;
- Automatic anchor signal discrimination may perform well for a majority of programmes, but may be tricked by similar signals or may not trigger at all, thus not giving 100% consistent results;
- File-based environments need a measurement paradigm that is applicable to 100% of the content and that delivers results that are 'good enough' for all programmes;
- Identifying an anchor signal needs input from an experienced operator or a discrimination algorithm; such an algorithm may be subject to the potential uncertainties listed above.

Anchor normalisation could offer better results on wide LRA material. It is however a task requiring expertise, and thus time and money, and if automatic discrimination is used, such an algorithm cannot be 100% reliable. Special measures need to be taken when anchor-adjusted content enters normalisation systems on file servers and thus needs bypassing the automatic processes in place. **As the biggest common denominator, R 128 recommends to measure the whole programme with all its elements instead of anchors, even with wide LRA material.**

## *5.2 Low Frequency Effects (LFE) Channel*

As noted in the description of ITU-R BS.1770 (see § 2), the LFE channel is currently **excluded** from the measurement. One of the reasons is the widespread uncertainty of consumers and audio engineers as well as equipment implementation differences regarding the alignment of this channel (+10 dB in-band gain). The omission of the LFE channel during the loudness measurement might cause its abuse. Further investigations of this matter and practical experience are needed to decide if and in which way the LFE signal might be included. One solution to completely avoid all potential issues with the LFE signal is not to use it at all if there is no need for extra headroom in the low bass region.

# 6. File-Based Production and Playout

As the broadcast world is changing to file-based workflows, it is vital that the loudness normalisation concept is also fully embraced there. The basic principle stays the same**: loudness normalisation** and **dynamic control** of the **audio signal** is recommended, especially for new content. Nevertheless, as **Metadata** is an integral part of file-based systems, solutions that rely more on Metadata are described as well (§ 7).

The origin of a broadcast file that contains audio signals can be via an ingest process, via transfer from an external server and from a file-based archive.

For existing programmes (archival content) there are basically four options to achieve loudness normalisation:

- Actually **changing** the the loudness level of **all audio files** to be 'on target'
- **Changing** the loudness level only '**on demand**'
- Using the result of a loudness level measurement to **adjust** the **playout level** without changing the original loudness level
- **Transporting** the **correct** loudness **Metadata** to the consumer where normalisation is performed

Which solution is actually chosen depends on factors like specific infrastructure, workflows, media asset management, availability of suitable equipment, financial resources, time etc.

At the very beginning of a file's life inside a facility, measurements have to be made, providing the values for **Programme Loudness Level**, **Loudness Range** and **Maximum True Peak Level** – the three characteristic audio parameters defined in EBU R 128 (Maximum Momentary and Maximum Short-term Loudness Level may be measured and stored too for very short content (<30 s, see § 7)). Depending on the results of that measurement and the subsequent method to achieve loudness normalisation and compliance with the acceptable Loudness Range, a processing scheme is executed, consisting of 'building blocks' or 'core tasks'. The workflow will now be examined in detail, with the help of generic flow-charts.

## 6.1 Building Blocks

**Programme Loudness Level Processing** (Figure 12)

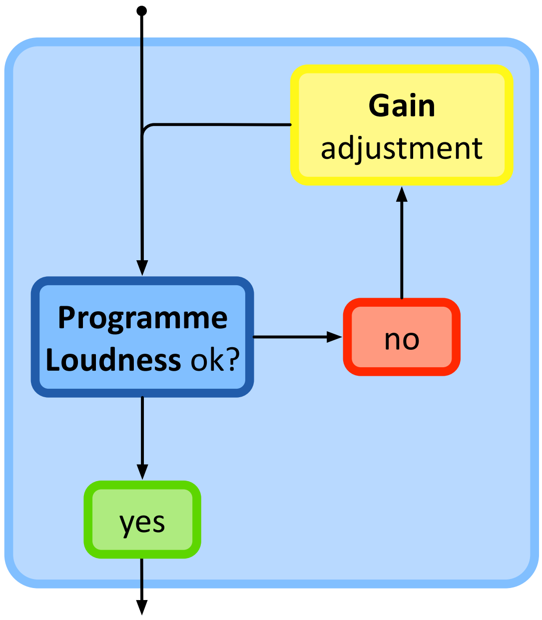


**Figure 12: Programme Loudness processing block**

**Loudness Range Processing** (Figure 13)

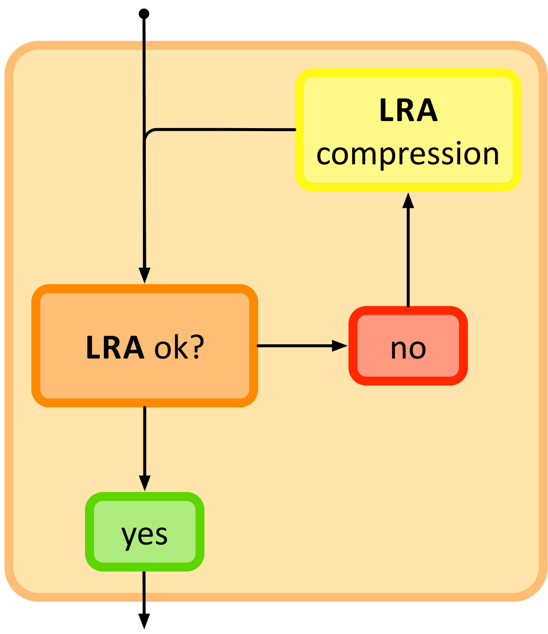


**Figure 13: Loudness Range processing block**

**Maximum True Peak Level Processing** (Figure 14)

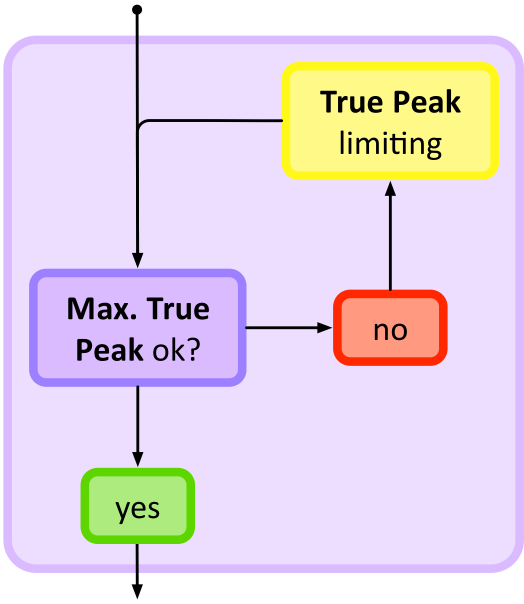


**Figure 14: Maximum True Peak Level processing block**

## 6.2 *Generic Loudness Levelling Strategies - Processing*

The three basic building blocks described above are at the core of any file quality control process regarding the technical parameters of its audio content. **At the beginning of any potential processing the three parameters Programme Loudness Level ($L_K$), Loudness Range (LRA) and Maximum True Peak Level (Max TP) are measured.** The result of this initial measurement determines the subsequent processing.

Several different scenarios are possible:

a) **All three parameters are OK.**


This is obviously the ideal outcome of the measurement: the **Programme Loudness Level** is **-23 LUFS**, the **Loudness Range** is within the specified limits of the broadcaster (depending on the genre and/or the distribution platform) and the **Maximum True Peak Level** is equal to or below the specified maximum value for the designated distribution system.

**b) The Programme Loudness Level is higher than -23 LUFS.**


A simple gain ranging (level reduction) operation solves that:

$$\textbf{Gain (dB)} = \mathbf{L_K Target - L_K measured}$$

(Example: the measured $L_K$ is -19.4 LUFS; Target Level is -23 LUFS; the necessary gain is [-23 - (-19.4) =] -3.6 dB. Max TP is naturally reduced by the same amount as $L_K$.)

**c) The Programme Loudness Level is lower than -23 LUFS.**


After applying a positive gain offset, the Maximum True Peak Level has to be recalculated (originally measured Max TP + gain offset = resulting Max TP) as it potentially lies above the permitted limit. If the new Max TP indeed exceeds the permitted limit, **True Peak Limiting** has to be performed according to the True Peak Processing Building Block. Another solution, which is applicable if such True Peak Limiting is not possible or wanted (or potentially too severe) is to leave $L_K$ at the original lower level and apply the **appropriate loudness Metadata setting** (lower than -23, reflecting the original loudness level). This requires that a fully functional system that supports and transports Metadata is in place (e.g. Dolby Digital, or MPEG-4).

For both scenarios **b** and **c** a simple **gain value** stored as Metadata may be used with potential subsequent limiting if Max TP is exceeded after a positive gain offset (scenario c). This gain value can control the playout level of the file so that -23 LUFS is reached.

**d) The Programme Loudness Level is lower than -23 LUFS and Loudness Range is wider than the internal tolerance for the genre or distribution channel.**


The Programme Loudness Level can be treated according to c above. Loudness Range is subject to processing (LRA Building Block) and thus potentially reduces Max TP. Although Max TP could have exceeded the permitted limit when applying a positive gain offset to $L_K$, Max TP Processing might not be necessary because of the LRA reduction. A calculation of Max TP during the LRA reduction process is therefore needed.

e) **Loudness Range is wider than the tolerance for the genre or distribution channel.**


As mentioned in § 4.4, a compressor with low threshold and a very moderate ratio can be used to narrow LRA (Loudness Range Building Block). For files, automatic processes with a 'target-LRA' are advantageous. Alternatively, the result of the LRA measurement might activate a dynamics compressor preset downstream with parameters similar to those listed in § 4.4. Max TP can only become lower, and so there is no potential for any True Peak Processing.

f) **The Maximum Permitted True Peak Level is exceeded.**


Exceeding the Max TP level of the respective distribution system incurs a risk of **distortion** further downstream (in a D-to-A-converter, sample-rate converter or bitrate reduction codec, for example). According to the Max TP Building Block True Peak Limiting is applied to lower Max TP. Whether there is a siginificant change to Programme Loudness as a result of this depends on the number and size of the peaks that are affected.

Any other combination of results of the initial measurement of $L_K$, LRA and Max TP are covered by processes already introduced in the scenarios above.

## 7. Metadata

As described in § 3.2, loudness normalisation can be either achieved through **normalisation of the audio signal** (the recommended method) or by **using Metadata** to store the actual loudness level. For the latter, the shift to Target Level can be performed either during the transfer of the audio file to the playout server, in the playout audio mixer, through choosing the appropriate preset of a downstream dynamics processor or directly at the consumer end with an adjustment of the playback level.

Metadata generally can be *active* (potentially changing the audio signal) or *descriptive* (providing information about the signal, such as format, copyright etc.). As a natural consequence of the work within PLOUD and the publication of EBU R 128 and its supporting documents, the three main parameters **Programme Loudness**, **Loudness Range** and **Maximum True Peak Level** shall form the core of loudness Metadata in audio files. Work is underway to include those parameters in the header (Broadcast Extension (BEXT) chunk) of the Broadcast Wave File (**BWF**) format (*for a detailed description of BWF, see [10], [11] and [12]*). Furthermore, the values for the **Maximum Momentary Loudness Level** as well as the **Maximum Short-term Loudness Level** shall be stored, as these parameters are helpful for controlling the dynamics of very short content (<30 s; see also § 10). Loudness Metadata is also intended to be included in the SMPTE dictionary with potential refinements like 'Loudness Profiles', addressing, for example, different processing presets of downstream loudness processors.

The Metadata parameters in existing systems that are of primary interest concerning loudness are:

- programme loudness
- dynamic range control words
- downmix coefficients

For example, in the Dolby AC-3 Metadata system, these parameters are called *dialnorm* (dialogue normalisation), *dynrng* (dynamic range) and *Centre/Surround Downmix Level*. The parameter *dialnorm* genuinely describes the loudness of an entire programme with all its elements such as voice, music or sound effects (also a music-only programme has a 'dialnorm' value). This may seem confusing; the reason is the focus of the Dolby system on normalisation according to the anchor signal dialogue.

## 7.1 Programme Loudness Metadata

Following the emphasis on **normalising the audio signal** in production to **-23 LUFS**, the relevant Metadata parameter shall naturally also be set to indicate **-23 LUFS**, provided the programme has been normalised to the Target Level. Consequently, after widespread normalisation of the source audio signals the Programme Loudness Metadata parameter will be **static**.

Exceptions where a different value than -23 may be used are:

- The programme does not fit into the window provided by -23 LUFS and -1 dBTP. This may occur mainly with very dynamic feature films and for broadcasters who want to transmit these programmes with such a large loudness-to-peak ratio;
- Legacy programming from the archive may not be able to be adjusted in time to fulfil the target level system of R 128;
- External live programmes may be provided with different loudness levels and Metadata.
- A fully functional system of providing and using Metadata over the whole signal chain is already in place. This implies faithful transportation of loudness Metadata to the consumer's home equipment.

In all these circumstances the **correct** Metadata value for Programme Loudness, measured with an 'EBU Mode' compliant meter, shall be set by all means. The distribution systems as well as Home Theatre Equipment handle this situation downstream (*see EBU Tech Doc 3344*).

## 7.2 Dynamic Range Control Metadata

Just as loudness normalisation can be performed at the source audio signal or via Metadata, the same applies to dynamic range processing. In the Metadata environment, dynamic range compression information is sent as part of the datastream in the form of *gain-words*. In the Home Theatre Equipment of the consumer, this information is applied to reduce the dynamic range of the signal, either by default or after user activation. Dynamic range control through the use of Metadata is not comparable with a sophisticated dynamics processor, but it provides a 'sticking plaster' for situations where the consumer wants a considerably lower dynamic range.

Referring again to the Dolby Digital system, there are *6 compression presets* that cause the encoder to generate different gain control words that are sent in the bitstream to the consumer's decoder: Film Standard, Film Light, Music Standard, Music Light, Speech and None. These presets result in more or less compression centred around the *dialnorm* value, one more reason to set this Metadata parameter correctly (*see Figure 15 for the compression curves around -23 LUFS*).

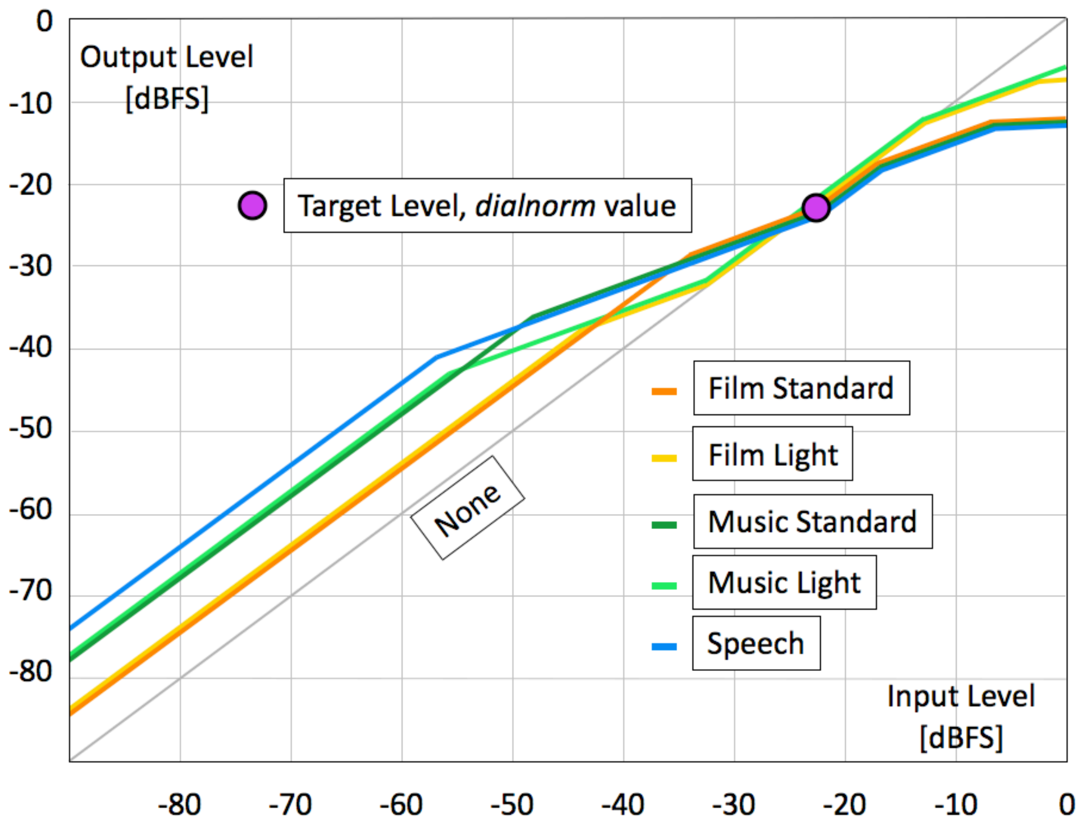


**Figure 15: Generic Dynamic Range compression curves of the AC-3 system**

Two compression profiles exist within Dolby Digital: *'Line mode'* and '*RF mode*'. For each, a separate compression preset can be chosen.

Within the system of **R 128** and its concept of normalising the audio signal to **-23 LUFS** as well as using the parameter Loudness Range to determine any potential processing, the preset '**None**' may be used. This may be applicable in particular for '*Line mode*' and also by default in '*RF mode*'.

The control of Loudness Range through actual processing of the audio signal at transmission is generally shifting the issue upstream. However, for specific programmes a broadcaster may choose a *gentle* profile for RF-mode systems (to avoid too active overload protection) while still choosing 'None' for Line-mode systems. Broadcasters that need other profiles than 'None' to support their internal workflow must be aware that this functionality may not always be implemented reliably in their listeners' equipment. **Manufacturers and distribution companies are advised to ensure that equipment is made in accordance with EBU Tech Doc 3344** ('*Distribution Guidelines*').

## 7.3 Downmix Coefficients

These Metadata parameters (again, as an example, here for Dolby Digital) are obviously only applicable to surround sound signals, controlling the gain (in dB) of the centre channel and the surround channels when mixed to Left front and Right front to derive a 2-channel-stereo signal. The loudness of a 2-ch-stereo signal, which is the result of an automatic downmix using Metadata, is dependent on:

- the actual downmix coefficients themselves (+3/+1.5/0/-1.5/-3/-4.5/-6/-∞)
- the programme content in the centre and surround channels and
- potential safety-limiting to avoid overload

Care should be taken to **avoid overload** of the downmixed signal. This can be achieved with a dynamics processor upstream. Static scaling (overall level reduction) should be avoided, as it systematically introduces loudness differences between the 2-channel stereo downmix and the original surround sound signal. Dynamic scaling may offer a solution.

The downmix coefficients possible within the Dolby-Digital system are governed by two downmix profiles. Initially, when there was only one profile, the parameters were coarser, with -3/-4.5/-6 dB for the Centre and -3/-6/-∞ dB for the Surround channels. Now, Extended Bitstream Information (Extended BSI) provides the finer intermediate steps listed above (in the first bullet point; also DVB TS 101 154 downmix coefficients offer the same resolution as the Dolby Digital Extended BSI). Broadcasters should be aware of the fact that not all reproduction equipment is able to deliver the intended downmix experience if Extended BSI is used, as legacy decoders may not be able to extract this information and would fall back to the fewer and coarser coefficients of profile 1.

In the case of missing or unreliable downmix Metadata, a good starting point is to look at the coefficients described in ITU-R BS.775-2 [13]:

L, R front: **0** dB

C, LS, RS: **-3** dB

It is also noted that the surround channels are weighted with +1.5 dB[5] during a loudness measurement according to ITU-R BS.1770. After an automatic downmix this weighting is not applied, as the result is only frontal 2-ch-stereo (Left and Right front). Programmes with a lot of surround content will consequently exhibit potentially larger variations of the loudness of the surround mix vs. the 2-ch-stereo downmix than programmes with more 'conservative' use of the surround channels.

[5] The **+1.5 dB weighting coefficient** for the surround signals in a **loudness measurement** according to ITU-R BS.1770 is not to be confused with the actual **+3 dB gain** for the surround signals in the **cinema**! In the cinema, the two individual surround channels are aligned 3 dB lower in level than the front channels, so that their combined level equals one front channel. The reason for this is compatibility with Mono-Surround movies (matrix-encoded 'Dolby Stereo' has only a (band-limited) mono-surround signal), where both surround channels would get the identical signal. For discrete multichannel audio mixes ('5.1' etc.) the surround signals in the final mix are therefore 3 dB 'hotter', as the mixing engineer compensates for the 3 dB lower alignment of the surround channels. If a cinema mix is broadcast, this 3 dB difference has to be compensated (other parameters like Loudness Range should be adapted as well).

Whereas the +3 dB gain for the surround channels has a purely technical reason, the +1.5 dB gain for the surround signals in the loudness measurement has psychoacoustic reasons. Humans perceive sounds coming from the back louder than frontal ones with the same sound pressure level. A measurement device does not have a brain and thus needs this gain factor.

In any case, there can be no guarantee that the Metadata supplied with an *external* file (or other media) are correct. Programme Loudness Metadata indicating -27 (the factory default for *dialnorm* in the Dolby-Digital system) or -31 (the lowest possible value in that system) are likely to raise special awareness, as chances are that Metadata have either not been looked at or been abused for the programme to appear (much) louder when replayed at the consumer's side.

It is therefore recommended to **discard loudness and dynamic range control Metadata** for external sources (except where the source can be fully trusted). Downmix coefficients may be passed through a fully functional Metadata system. An entire measurement process of the three main audio parameters needs to be conducted afresh. Only this will ensure the correct subsequent processing. For internal purposes, Metadata can be better controlled.

# 8. Alignment of Signals in the Light of Loudness Normalisation

## *8.1 Alignment Signal and Level*

An **Alignment Signal** in broadcasting consists of a sine-wave signal at a frequency of 1 kHz, which is used to technically align a sound-programme connection. In digital systems the **level** of such an Alignment Signal is **18 dB** below the maximum coding level, irrespective of the total number of bits available (**-18 dBFS**). The switch to loudness normalisation does **NOT** change this approach, as alignment does not imply a mandatory connection to loudness metering or measurement.

Therefore, **alignment for sound-programme exchange can be performed as usual, with a sine-wave signal of 1 kHz at a level of -18 dBFS**.

**The alignment level for sound-programme exchange does not need to change. Use a 1 kHz sine wave at -18 dBFS as usual.**

This is specified in EBU Recommendation R 68 [14]. In the same document the "Permitted Maximum Level" is still mentioned, as defined in ITU-R Recommendation BS.645-2 [15]; With the change to the "Maximum Permitted True Peak Level" (-1 dBTP for generic PCM productions) being different than the recommended -9 dBFS in ITU-R BS.645 (because of the QPPM-scenario becoming obsolete), the relevant sections of EBU R 68 – 2000 and ITU-R BS.645 (as well as documents that refer to the definition of "Permitted Maximum Level" within these recommendations) potentially need to be revised.

The alignment level of **-18 dBFS** (1 kHz tone) will read as **-18 LUFS** on a loudness meter with absolute scale (or +5 LU on the relative EBU-mode scale), provided that the 1 kHz tone is present (in phase) on both the left and right channel of a stereo or surround sound signal. If the 1 kHz tone is used only in a single front channel, the loudness meter will read -21 LUFS (or +2 LU on the relative scale).

**A stereo 1 kHz sine wave at -18 dBFS reads as -18 LUFS absolute (+5 LU relative) on an EBU mode loudness meter.**

## 8.2 *Listening Level*

A different topic is the **Listening Level** in an audio reproduction system. In the relevant document, EBU Tech Doc 3276-E '*Listening conditions for the assessment of sound programme material*' (and Supplement 1, extending it to Multichannel Sound), the following formulas are used to adjust the level of one loudspeaker [16]:

(1) $L_{LISTref}$ = 85 – 10log2 $dB_A$ (for 2-channel stereo)

(2) $L_{LISTref}$ = 96 $dB_C$, referenced to digital full scale signal level (for multichannel audio up to 5.1)

To achieve this, a signal consisting of noise of equal energy per octave and covering either the whole frequency range (equation (1)) or the frequency range from **500 Hz** to **2 kHz** (equation (2)) should be employed. Measurements should actually be made at a mean signal level equal to the **alignment level**, which is defined here as 18 dB below digital full scale. Under these conditions, the loudspeaker gains should be adjusted to achieve a **Reference Listening Level** ($L_{LISTref}$) of 85 - 3 = **82** $dB_A$ Sound Pressure Level (SPL) per loudspeaker for 2-channel stereo systems and 96 – 18 = **78** $dB_C$ SPL for multichannel systems. The measurements should be made at the reference listening position using an A-weighted slow response sound level meter for 2-channel stereo and a C-weighted one for multichannel audio.

This is arguably a bit confusing, with different numbers with different noise signals and different frequency weighting of the sound level meter. But these differences actually compensate in a way and result in a similar listening level for both 2-channel stereo and multichannel systems.

***To summarize:***

| | | |
|---|---|---|
| For 2-channel stereo: | $L_{LISTref}$ = **82** $dB_A$ SPL per loudspeaker | (using 20 Hz – 20 kHz noise of equal energy per octave at -18 dBFS rms) |
| For 5.1 MCA: | $L_{LISTref}$ = **78** $dB_C$ SPL per loudspeaker | (using 500Hz – 2 kHz noise of equal energy per octave at -18 dBFS rms) |

It is anticipated that the average level of sound-programmes will be *lower* once EBU R 128 is put into practice. The decrease in level may be in the order of up to 3 LU (in extreme cases even more). This makes a corresponding **increase** in the monitor level of the reproduction system seem likely. As indicated above, the *Alignment Level* does not have to change accordingly, as the alignment procedure is still valid to ensure a reasonable gain structure as well as a high signal-to-noise ratio in the reproduction chain. If there is widespread agreement in the future on an increased level to the monitors due to the lower average level, the relevant documents will be examined again.

**As -23 LUFS is about 3 LU lower than today's average programme level, one might consider raising the level sent to the monitors accordingly.**

## 9. Implementation and Migration

It is clear that such a fundamental change in the way audio signals are measured, metered and treated, and that affects all stages of audio production, distribution, archiving and transmission, is not done overnight with the flick of a switch. Every broadcaster and audio facility must find its individual way to perform this change, to install the appropriate equipment, train the staff and get on the road to loudness paradise! Nevertheless, a few constants can be stated that will be applicable for everyone. They are presented in the following sections.

### *9.1 Generic Migration and Implementation Advice*

- Establish an **internal loudness group** to discuss basic implications and a strategy to convince management, programme makers and your colleagues.
- **Start now** – don't wait until everything is in place and all the others have done it and don't try to be perfect in the very beginning.
- Before you can do anything, management has to agree to this change and all its consequences. Get a **written agreement** or 'call for action' from the general director.
- Provide **loudness meters** to your key production personnel. Let them play with them, gain first experiences and learn the advantages and liberations of the loudness paradigm so they can be opinion leaders for their colleagues.
- **Survey the market** regarding loudness metering and loudness management to determine what is best suited for your environments.
- Determine the **key areas** in which loudness work should start. Potential candidates are: production studios, post-production suites, OB vans, QC (Quality Control) department.
- Be aware that you will encounter obstacles ("It has always been that way", "It has never been that way", "Who are you to say we should do it that way"). **Patience** and demonstrating practical examples will pay off. Become your facility's Zen-master of loudness normalisation ("restraint – simplicity – naturalness").
- **Allow** everybody **time to adapt**. Although the audience has been waiting for a solution for decades, don't create more problems by trying to do too much too quickly.
- Solutions for **file-based workflows** are still rare (February 2011). Keep an eye on the market and demand solutions from vendors.
- Use this fundamental change as an opportunity for a general **discussion** about **audio quality** and the development of a '**corporate sound**', which includes, for example, speech intelligibility, the balance of speech vs. music and, of course, loudness normalisation of programmes.
- **Use and trust your ears**! They are the best loudness meters. Smile when you watch your colleagues working with the loudness paradigm as if nothing else ever existed.

### *9.2 10 Points of Action for Migration and Implementation*

- **Establish an internal loudness group.**
- **Don't wait, start now.**
- **Get a written agreement from management.**
- **Give your key personnel loudness meters.**
- **Survey the market for loudness equipment.**
- **Choose the key areas to be the first to change.**
- **There will be obstacles. Be patient, allow time.**
- **Use the momentum to discuss audio quality.**
- **Use and trust your ears.**
- **Become your facility's Zen master of loudness.**

## 10. Genre-Specific Issues

The concept of **EBU R 128** centres on the loudness normalisation of each programme to one single Target Level (**-23 LUFS**). There are two reasons why this cannot be a perfect solution:

- No objective loudness measurement can ever be perfect
- There will always be individual preferences

Thus, a perfect solution is *generally* not possible as it differs from person to person. Within the scope of EBU R 128 it is vital to understand that it is not intended to achieve a loudness balance based on the real sound pressure level of a specific audio signal, but instead provide a satisfactory listening experience for a diverse mix of genres for the majority of listeners.

This results, for example, in a Schubert string quartet having the very same integrated loudness level as a Mahler symphony, namely -23 LUFS. While this does not reflect reality, it makes these items fit into a wide array of adjacent programming, and that is the intention of **advocating one single number**.

As this document shall serve as a pool of experiences, one might be tempted to consider refining this paradigm, once loudness normalisation becomes widespread. But listeners do accept if the loudness level of programmes lies within a so-called '**comfort zone**' of about **8-9 LU**, whereas the distribution is asymmetric (for example, +3 LU/-5 LU). In cases, where the objective loudness algorithm does not always provide a perfect result, the programme will most certainly still lie within this comfort zone. Broadcasters should also bear in mind that the audience can still adjust the loudness level with their remote control, to accommodate likes and dislikes.

The EBU encourages the **normalisation to one single target level** despite a potential refinement for individual genres. Allowing too many variations (or even only a few) may challenge the system of equal average loudness from the onset. Naturally, the fear is that variations would be biased to the louder side.

Exhibiting a Programme Loudness Level **lower** than the Target Level is a slightly different topic. As a 'case study', two genres are now investigated, where a specific treatment (also for the Maximum Loudness Level) may be appropriate under certain circumstances; **commercials** (advertisements) and **trailers**, and **music** programmes.

## *10.1 Commercials (Advertisements) and Trailers*

This type of programme is arguably the most frequently mentioned one as far as **listener annoyance** is concerned, and thus is mainly responsible for the loudness problems encountered today. In the UK (BCAP rules – Broadcast Committee of Advertising Practice) and the USA (CALM Act – Commercial Advertisement Loudness Mitigation) even *legislation* has recently been put into place to tame this genre. It is certainly vital that the system of loudness normalisation based on **EBU R 128** provides an effective toolset for this task – abuse shall be prevented. To control the dynamics of a commercial in a loudness normalised world where there exists the danger of suddenly too high loudness differences (overly loud 'pay-off' after a longer period of low-level signals just above the gate threshold), the parameter Loudness Range (LRA) is not suited, as the calculation is based on the short-term loudness values (3 s interval). Therefore, for very short items there are too few data points to derive a meaningful number for LRA. The Loudness Range parameter is not to blame for this fact, as it was never intended for this purpose.

An alternative can be found in using the **Maximum Momentary Loudness Level** (Max ML – 400 ms) and/or the **Maximum Short-term Loudness Level** (Max SL – 3 s). Especially for very short items (<30 s), these parameters can be effectively used to limit loudness peaks. First experience of PLOUD members has pointed to a value around **+8 LU** (-15 LUFS) as a possible limit for Max ML and **+3 LU** (-20 LUFS) for Max SL. In any case, both parameters (Max ML and Max SL) are part of the suggested extension of the Metadata of the Broadcast Wave File Format (BWF). EBU members are encouraged to use individual limits of Max ML or Max SL for short items and report their findings.

**A limit for Maximum Momentary and Maximum Short-term Loudness may be used to prevent abuse for very short items (< 30 s).**

For programmes of this genre that consist of only background or creatively wanted low-level sounds, a loudness level **lower** than Target Level may be used. This is in line with past and current practice to limit the maximum peak level (now: loudness level), but not enforce all content to sit at that maximum. Deliberate low-level audio provides **contrast**, and this is one of the most fundamental creative tools in every art form. The short duration of the commercials, advertisements or trailers that might use this dramaturgical tool effectively is not likely to cause any influence on the daily long-term average loudness level of the station.

**Programmes destined for playout lower than Target Level need special attention to ensure they pass automatic normalisation processes unharmed. They should really be the exception, not the rule.**

Ultimately, the responsibility for all these cases and decisions lies with the producer, director or other creative personnel, respectively.

## *10.2 Music*

The experience of passionate music listeners suggests that certain programmes that contain mostly music, either with a wide loudness range like classical music or with a higher degree of dynamic compression as an artistic property like a rock concert, have the tendency to be listened to with a higher loudness level (up to **+2-3 LU** on average) than other genres. Reasons for that might be the significantly high potential sound pressure level in reality (fortissimo of a symphony orchestra, rock band with powerful public address system) and the fact that for music there do not exist 'foreground sounds' vs. 'background sounds' – everything is in the foreground.

But as mentioned above, a potential differentiation of the target level for these programmes may cause more harm in opening a backdoor to being again louder than the rest instead of improving the situation significantly. Based on the same reasoning as for commercials, advertisements and trailers, normalisation to a different (= higher) target level is **discouraged**. The audience can still use their remote control to adjust (increase) the loudness level in their reproduction environment to their taste. Adjacent programmes like commercials or trailers will consequently be shifted too. It is anticipated that this should not push those programmes out of the comfort zone.

# 11. References


[1] EBU Technical Recommendation R 128 '*Loudness normalisation and permitted maximum level of audio signals*' (2010)

[2] ITU-R BS.1770 '*Algorithms to measure audio programme loudness and true-peak audio level*' (2006-2007)

[3] EBU Tech Doc 3341 '*Loudness Metering: 'EBU Mode' metering to supplement loudness normalisation in accordance with EBU R 128*' (2010)

[4] EBU Tech Doc 3342 '*Loudness Range: A descriptor to supplement loudness normalisation in accordance with EBU R 128*' (2010)

[5] EBU Tech Doc 3344 '*Practical Guidelines for Distribution Systems in accordance with EBU R 128*' (2011)

[6] ISO 80000-8: '*Quantities and Units – Part 8: Acoustics*'

[7] Grimm E., Skovenborg E. & Spikofski G. '*Determining an Optimal Gated Loudness Measurement for TV Sound Normalization*', AES Convention Paper N° 8154, 128th AES Convention, May 2010

[8] ITU-R BS.1864 '*Operational practices for loudness in the international exchange of digital television programmes*' (2010)

[9] Lund, Th. '*Stop counting samples*', AES paper N° 6972, 121st AES Convention, October 2006

[10] EBU Technical Recommendation R 85: '*Use of the Broadcast Wave Format for the Exchange of Audio Data Files*' (2004)

[11] EBU Technical Recommendation R 111: '*Multi-channel Use of the BWF Audio File Format (MBWF)*' (2007)

[12] EBU Tech Doc 3306: '*MBWF/RF64: An extended File Format for Audio*' (2009)

[13] ITU-R BS.775-2 '*Multichannel stereophonic sound system with and without accompanying picture*' (2006)

[14] EBU Technical Recommendation R 68: '*Alignment level in digital audio production equipment and in digital audio recorders*' (revision 2000)

[15] ITU-R BS.645-2 '*Test signals and metering to be used on international sound programme connections*' (1992)

[16] EBU Tech Doc 3276-E (+ supplement 1) '*Listening conditions for the assessment of sound programme material*' (1998, 2004 – supplement 1)