

amagi THUNDERSTORM



Audio Normalization

Overview

Audio Normalization is the process of automatically adjusting the loudness of every piece of content – ads, promos, fillers, and stitched outputs – so that all audio plays back at a consistent volume, regardless of how loud or quiet the original source file was mixed.

This eliminates the most common viewer complaint in ad-supported streaming: sudden volume jumps between programme content and advertising.

Why It Matters

Benefit	Impact
Consistent viewer experience	Ads blend smoothly into programming, no startling volume jumps.
Higher engagement & retention	Viewers don't reach for the remote or skip ads due to discomfort.
Regulatory compliance	Meets US CALM Act, EU EBU R128, and equivalent international loudness regulations out of the box.
Advertiser parity	Every advertiser gets equal sonic presence, no creative is unfairly buried or unfairly louder than others.
Device safety	Prevents clipping, crackling, and distortion on TVs, mobiles, and in-car audio systems.

Technical Standard

Amagi Thunderstorm measures and corrects loudness using the international standard **ITU-R BS.1770**, the basis for both:

- **ATSC A/85 (CALM Act)** : United States broadcast standard
- **EBU R128** : European broadcast and OTT standard

Loudness is measured in **LUFS** (*Loudness Units relative to Full Scale*) a perceptual unit that reflects how loud the human ear actually perceives the audio, rather than just the raw signal amplitude.

Default Normalization Targets

Every audio stream processed by the Amagi Thunderstorm is normalized to the following targets:

Parameter	Target Value	Purpose
Integrated Loudness	-24 LUFS	Average perceived loudness across the clip (US ATSC A/85 / CALM Act compliant)
True Peak	-2 dBTP	Maximum allowable peak; provides 2 dB of headroom to prevent clipping and codec distortion on downstream devices
Loudness Range (LRA)	7 LU	Controlled dynamic range – preserves natural dynamics while keeping dialogue intelligible

These targets are configurable per customer or per market. For example:

- EBU R128 delivery: -23 LUFS / -1 dBTP
- OTT / streaming delivery (AES TD1004): -18 LUFS or -16 LUFS
- Custom targets can be set at the workflow level on request.

Coverage Across Output Formats

The same normalization is applied uniformly across every delivery path Amagi

Thunderstorm supports, ensuring consistency no matter which format the

end-viewer's player requests:

- HLS (Apple HTTP Live Streaming)
- DASH / CMAF (chunked streaming, low-latency live)
- Remuxed outputs
- VOD Just-In-Time stitched outputs where ad and filler segments are normalized to match each other, ensuring no audible seam at the splice point.

Edge-Case Handling

Amagi Thunderstorm is designed to be robust against common challenges observed

in real-world ad creatives:

- **Source files with no audio track** - A correctly formatted silent audio stream is automatically generated, so downstream players never fail due to a missing audio channel. Customers may alternatively choose to reject silent creatives at ingestion this is a configuration option.

- **Mono, stereo, and multi-channel inputs** - All are accepted and standardised to the channel layout defined by the customer profile (stereo by default).
- **Mixed sample rates** - All output is normalized to a single, profile-defined sample rate, ensuring consistency across an ad pod.

Monitoring (Roadmap)

A per-segment loudness telemetry is added to the transcoding pipeline. For every asset that is processed, following is recorded and published

Measurement	Descript
Input loudness (LUFS)	Source mezzanine measurement
Output loudness (LUFS)	Delivered measurement
True-peak level (dBTP)	Headroom against clipping and distortion
Correction offset (LU)	How far the output landed from the target