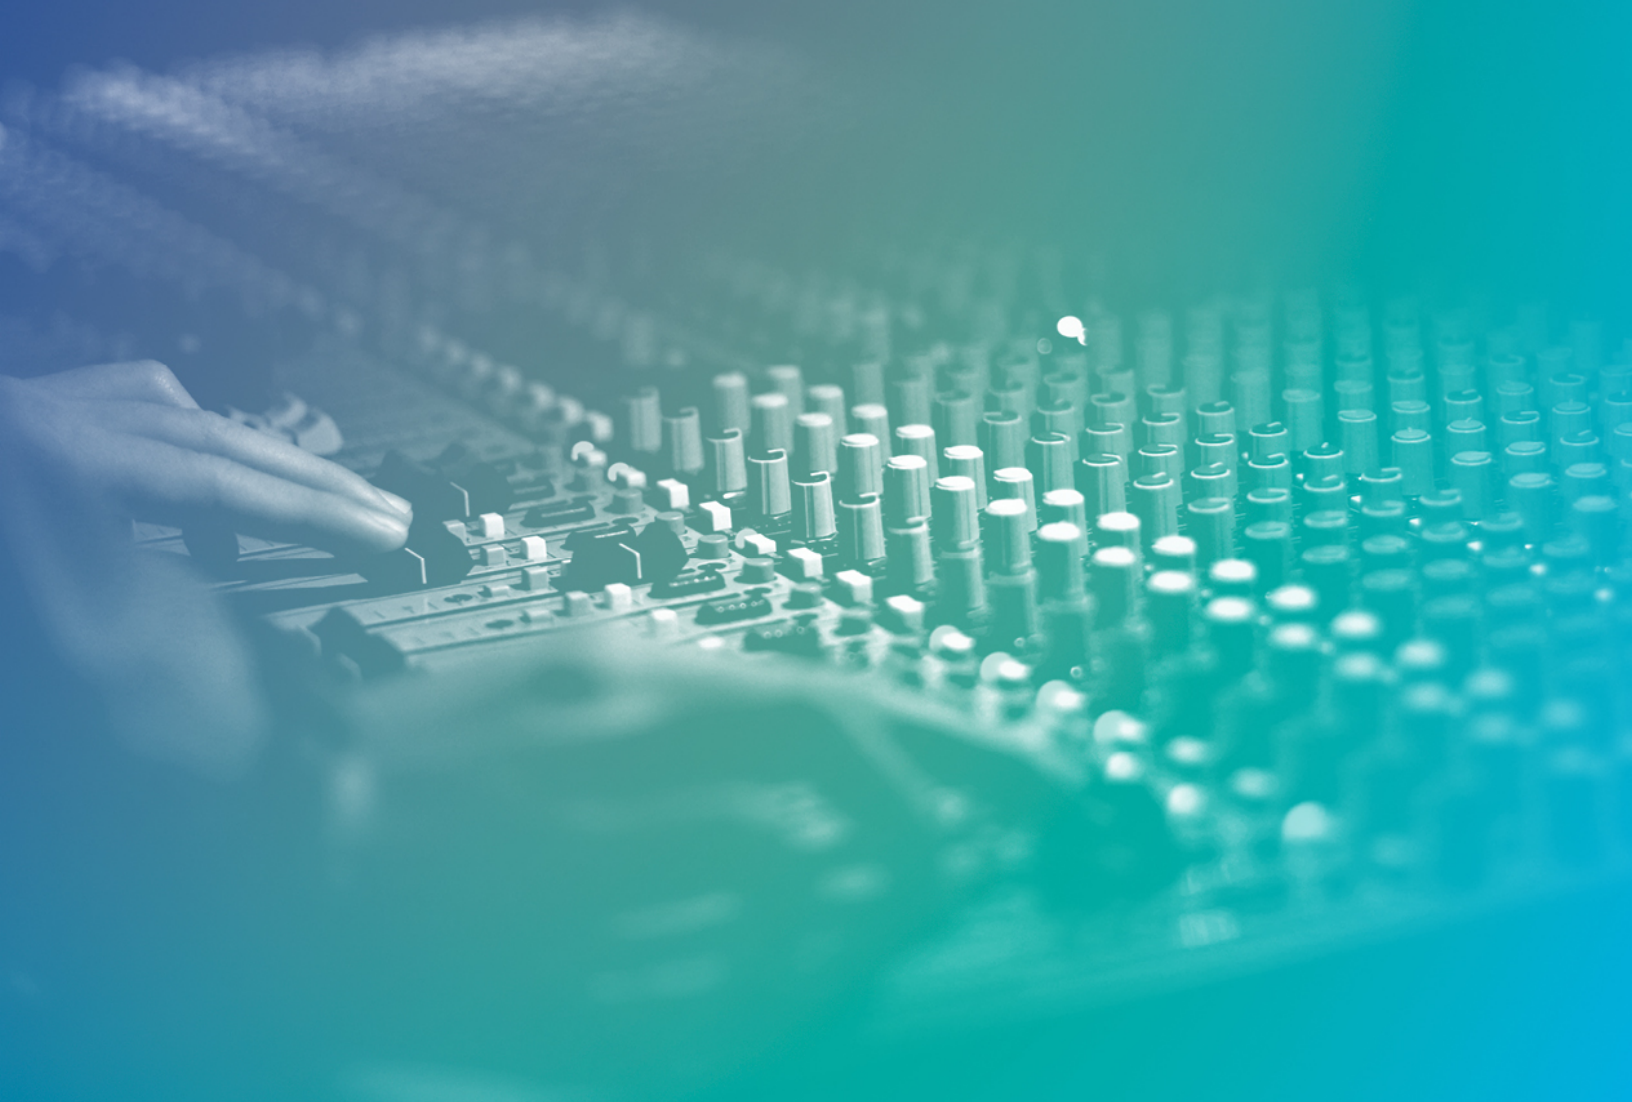




ATSC 3.0 Dolby Audio Handbook

February 2020

This handbook provides ATSC 3.0 broadcasters in North America with the documentation, support, and resources required to move from AC-3 to AC-4 audio in the transition from ATSC 1.0 to ATSC 3.0, with a specific focus on the practical and technical considerations integrating AC-4 into ATSC 3.0 workflows.



Contents

Introduction	2
--------------	---

Enabling ATSC 1.0 audio use cases for ATSC 3.0 using Dolby AC-4	3
• Updating workflows to enable ATSC 1.0 audio use cases with AC-4	4

Adding Next Generation Audio experiences	6
• Dolby Atmos	6
• Efficient multi-language delivery using audio presentations	6
• Metadata for Next Generation Audio	7
• AC-4 encoding products for NGA	10

MVPD contribution	11
• Sample workflow	12

Additional information	13
• AC-4 decoding in ATSC 3.0 TVs and STBs	13
• Emergency audio	14
• Ad insertion	14

Getting started	16
• Resources	16

Abbreviations	17
---------------	----

Introduction

The Advanced Television Systems Committee (ATSC), through a cooperative effort by over 125 member organizations from the broadcast, consumer electronics, cable, satellite, motion picture, professional broadcast equipment, computer, and integrated circuit industries, has developed the ATSC 3.0 television standard.

Terrestrial television continues to evolve, utilizing the latest technologies to reach viewers and offer a wide variety of program choices. Next Generation Audio (NGA) is a key part of ATSC 3.0, enabling broadcasters to deliver new experiences to the home.

ATSC 3.0-enabled receiving devices with new audio features through AC-4 decoding make it possible for broadcasters to offer completely new experiences.

Some of these new audio experiences will require broadcasters to modify their workflow's upstream of audio encoding. Broadcasters not looking to alter their current workflows will be able to continue offering the same audio experience they are delivering today in ATSC 1.0 with AC-3 through ATSC 3.0 with AC-4. Those broadcasters will benefit from some of the new features of AC-4 that require no modifications to existing workflows.



Enabling ATSC 1.0 audio use cases for ATSC 3.0 using Dolby AC-4

ATSC 3.0-specified AC-4 audio compression is two generations ahead of AC-3, the audio codec for ATSC 1.0. But AC-4 can still be used to deliver existing use cases from ATSC 1.0. Broadcasters looking to maintain their current ATSC 1.0 audio workflows indefinitely, or for an interim period of time, will not require a wholesale change of equipment and infrastructure to transition to AC-4.

DIALOG ENHANCEMENT

AC-4 allows viewers to tailor the dialog level to suit their individual preferences. These flexible mechanisms work with both legacy content that contains dialog mixed into the main audio and new content where a separate dialog track is available to the Dolby AC-4 encoder. The term 'Dialog Enhancement' is the technical name of this feature. For example, some AC-4 encoders use this name under the audio setting for AC-4. However, the ATSC has standardized on the term 'Voice Boost' to describe this feature for consumer facing applications. For example, this term would be used in a television's user interface to let the end user know that the Dialog Enhancement feature is available.

All AC-4 encoders perform this analysis and create a parametric description of the dialog. All AC-4 decoders in televisions and STBs (set top boxes) use this parametric description to perform Dialog Enhancement. There is no action needed by the broadcaster from an operational side to enable or activate this feature.

The AC-4 decoder implementation deliverables available through Dolby Laboratories to SoC (System on Chip) vendors, and television and STB manufacturers support Dialog Enhancement in all instances. However, the way this feature is exposed to the end user through the device's user interface is outside the control of Dolby Laboratories.

DATA RATE EFFICIENCY

AC-4 provides significant compression efficiency improvement over AC-3. For example, AC-3 is commonly used at a data rate of 384 kbps for delivering a 5.1-ch mix. With AC-4, broadcasters are able to use a data rate of 128 kbps to deliver the same 5.1-channel audio with the same quality. In order to provide margin for operations that might require downstream transcoding or other operations, the table below shows the recommended minimum data rates for stereo and 5.1-ch configurations.

Channel Config	Recommended AC-4 Data Rate	AC-3 Equivalent	Key
Stereo (2.0) CM	64 kbps	192 kbps	CM – Complete Main AD – Audio Description
5.1 CM	144 kbps	448 kbps	
Stereo CM + Stereo CM/AD	128 kbps	384 kbps (192+192)	
5.1 CM + Stereo CM/AD	208 kbps	640 kbps (448+192)	

INTEGRATED LOUDNESS MANAGEMENT USING RTLL (REAL TIME LOUDNESS LEVELER)

The AC-4 encoder has an integrated, non-destructive loudness leveling processor known as the Real Time Loudness Leveler. RTLL normalizes the incoming audio signal prior to encoding in AC-4. There is no action needed by the broadcaster from an operational side to enable or activate this feature.





If the incoming audio presented to the Dolby AC-4 encoder has previously been produced or adjusted to a target loudness level by a trusted device, this can be signaled to the encoder using metadata. In this case, the RTLL will be automatically disabled so that the audio is delivered without further adjustment, maximizing quality and preserving the original creative intention.

Updating workflows to enable ATSC 1.0 audio use cases with AC-4

For those broadcasters looking to maintain parity with AC-3 audio, the following equipment changes will need to be made when implementing AC-4 in an ATSC 3.0 platform:

AC-4 AUDIO ENCODER

Some HEVC video encoders offer integrated AC-4 audio encoding (internally, in parallel with the video encoding). At the time of this publication, the following video encoder manufacturers offer support for integrated AC-4 encoding:

			
harmonicinc.com	dsbroadcast.com	ateme.com	synamedia.com

Alternatively, broadcasters can encode AC-4 before their transmission feed to the HEVC video encoder using standalone audio encoders. The following audio encoder manufacturer currently offers support for AC-4 encoding:



AC-4-ENABLED DASH PACKAGER

ATSC 3.0 introduces a new technology for the Transport Layer, moving away from MPEG-2 Transport Streams and adopting MPEG DASH within ROUTE. Consequently, it is necessary for broadcasters to ensure that their MPEG DASH packagers, whether integrated within the video encoder or standalone, are conformant to the ATSC 3.0 specification and are capable of properly handling AC-4 audio streams.

The following manufacturers currently offer support for AC-4 DASH packaging:



AC-4 ANALYSIS AND MONITORING TOOLS

StreamXpert from DekTec ([dektec.com](https://www.dektec.com)) now supports monitoring the ATSC-3 Transport Layer including analysis and decoding of AC-4. A tool known and understood by many broadcasters, StreamXpert will display the AC-4 stream configuration, allow the user to access specific metadata parameters, and monitor selected audio.

The BGD4100 4K UHD/HD IRD from DS Broadcast ([dsbroadcast.com](https://www.dsbroadcast.com)) supports off-air reception of ATSC 3.0 signals, display of statistics, HEVC video decoding, and passthrough and decoding of AC-4 audio.



Adding Next Generation Audio experiences

To take advantage of some of the new audio features offered in ATSC 3.0, broadcasters may need to adapt parts of their workflow prior to the transmission encoder.

Dolby Atmos

In ATSC 1.0, soundtracks are limited to the traditional 5.1 channel surround configuration supported by AC-3. In ATSC 3.0, broadcasters can offer a more immersive experience through Dolby Atmos.

A common configuration for delivering Dolby Atmos in Broadcast is 5.1.4. This configuration adds a third dimension of overhead audio channels to the traditional two-dimensional 5.1 channel configuration.

A 5.1-channel mix is carried as 6 channels of PCM audio.

- Left Channel
- Right Channel
- Center Channel
- LFE Channel
- Left Surround Channel
- Right Surround Channel

In contrast, a 5.1.4-channel mix requires 10 channels of PCM audio.

- Left Channel
- Right Channel
- Center Channel
- LFE Channel
- Left Surround Channel
- Right Surround Channel
- **Top Front Left Channel**
- **Top Front Right Channel**
- **Top Back Left Channel**
- **Top Back Right Channel**

Broadcasters looking to offer Dolby Atmos via ATSC 3.0 need to be aware of the requirement for additional audio channels.

Efficient multi-language delivery using audio presentations

Traditional broadcast audio provides a single full mix of the program, known as CM (Complete Main) mix. A CM mix consists of dialog, music, and effects elements. AC-4 supports delivery of CM content for stereo, 5.1-ch and Dolby Atmos content.

In addition to working with CM mixes, AC-4 is able to deliver a single common M&E (music and effects) element together with separate D (dialog) elements, with each dialog element

selected and mixed with the M&E element in the AC-4 decoder. The AC-4 bitstream design allows a set of audio elements to be combined into multiple audio presentations, enabling multiple versions of the audio to be delivered in a single bitstream in a bandwidth-efficient manner.

Broadcasters looking to make use of compositional channel configurations need to produce content in this channel configuration, as well as preserve this channel configuration from the point of content creation all the way into the AC-4 encoder in transmission. Broadcasters may also receive content already in compositional channel configuration, in which case they will need to preserve this channel configuration through to the AC-4 encoder in transmission.

AC-4 allows a set of audio elements to be combined into multiple audio presentations, enabling multiple versions of the audio to be delivered in a single bitstream.

Broadcasters looking to offer Dolby Atmos and/or compositional channel configurations through AC-4 must be able to deliver audio metadata alongside the audio stream feeding the AC-4 encoder. This metadata must accurately represent the channel configuration(s) and all the other metadata parameters required by the AC-4 decoder for correct playback.

Metadata for Next Generation Audio

For correct audio playback, all AC-4 audio streams must carry metadata that correctly represents the characteristics of the audio content applied to the input of the encoder.

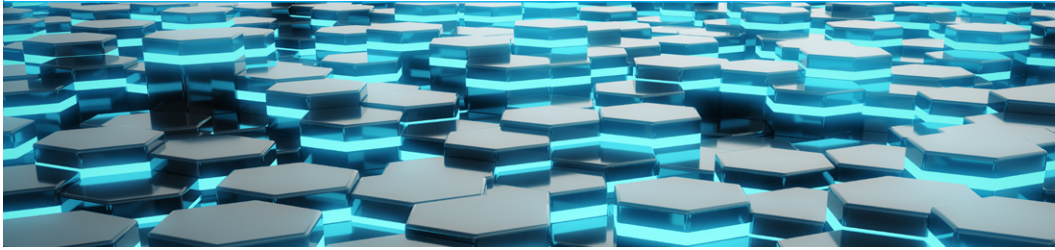
This metadata can be created at any stage prior to the AC-4 encoder, or within the AC-4 encoder itself. In either case, it is necessary that the metadata accurately reflects the characteristics of the audio content, such as loudness, channel configuration, or the various presentation selections when available.

GENERATING AUDIO METADATA DURING PRODUCTION

Audio metadata for Next Generation Audio can be created during audio production. This metadata can be carried from production to the AC-4 encoder using PMD (Professional Metadata). PMD is a new professional audio metadata format published in SMPTE RDD 49. PMD flows alongside the audio, and when using PCM audio in SDI or MADI, PMD occupies a channel. The carriage of PMD over SMPTE ST 337/AES3 is standardized in SMPTE ST 2109.

The AC-4 emission encoder will apply the metadata during encoding, or convert it to AC-4 bitstream metadata, as appropriate. PMD is an open format and is provided to the industry at no cost and with no requirement for a license. More information can be found at:

<https://developer.dolby.com/globalassets/professional/documents/pmd-application-guide.pdf>



GENERATING DEFAULT AUDIO METADATA AT THE ENCODER

LOUDNESS METADATA

If the input to the AC-4 audio encoder is PCM audio without metadata, the AC-4 encoder will generate AC-4 bitstream metadata using pre-selected values in the encoder. As with AC-3, it is important to ensure that the preset values match the incoming audio. For example, if the dialnorm value of an AC-4 encoder is set to '-24', then PCM content without metadata should be normalized to match that value prior to the encoder (although the RTLL process built into AC-4 will ensure correct playback loudness if there is a mismatch).

CHANNEL CONFIGURATION METADATA

In some ATSC 1.0 workflows, AC-3 encoders are always configured to output a 5.1-ch AC-3 bitstream. This ensures that when the AC-3 encoder is fed with a 5.1-ch PCM audio mix (without metadata), the channel configuration matches the AC-3 metadata. When stereo PCM audio is fed into an AC-3 encoder that is configured to output a 5.1-ch bitstream, the stereo audio is encoded into the L and R channels of the 5.1-ch AC-3 bitstream. Although not a recommended practice, this method enables some broadcasters to deliver stereo and 5.1-ch audio programming using a simple and relatively functional workflow that does not require changing the channel configuration of the AC-3 encoder dynamically. This same method would produce a similar result when using AC-4.

However, Dolby Atmos and/or compositional channel configurations with multiple presentations are not supported by this workflow, and attempting to use this approach could result in important audio elements, such as alternate dialog, appearing from inappropriate locations. The recommended practice is to always configure the AC-4 encoder using metadata that matches the content.

ENABLING NGA USING AC-4 IN CONSTRAINED WORKFLOWS

For existing SDI-based workflows, the number of provisioned channels available, as well as the ability to carry metadata along with the audio, can present challenges to enabling a full NGA experience. As a workaround for these limitations, it is possible to place the AC-4 encoder near the start of the audio workflow and pass the encoded AC-4 bitstream through the rest of the audio path using SMPTE ST 337 formatting.

Typically, when passing encoded bitstreams through an audio workflow, applications such as editing, ad insertion, and international turnarounds become challenging to implement, as the switching points at the end of video frames do not align with the end of audio frames.

Because AC-4 can be aligned with video frames, it is compatible with switching and insertion at video frame boundaries.

To solve these challenges, the Dolby AC-4 encoder features an optional video reference input to align the audio and video frames. The encoded audio frame rate can therefore be set to match the video frame rate and, as a result, the boundaries of the audio frames can be precisely aligned with the boundaries of the video frames. Because AC-4 can be aligned with video frames, it is compatible with switching and insertion at video frame boundaries.

For example, a broadcaster produces content, such as local news, as a stereo CM, and delivers this content as PCM without metadata into the transmission AC-4 encoder. The transmission encoder metadata preset is selected to match the incoming signal.

When the broadcaster switches the input source from the current stereo PCM to an external Outside Broadcast (OB) signal already encoded in AC-4, the switch between sources can be done at the video frame boundary without losing any audio frames or introducing A/V sync errors.

In this scenario, the broadcaster uses a simple *metadata preset workflow* for the stereo CM content, while continuing to offer advanced content with more complex metadata.

Note that for more complex audio workflows — editing elements within the audio during post production, or inserting audio watermarks — the AC-4 bitstream would have to be decoded and re-encoded, and metadata must be passed from the AC-4 decoder to AC-4 encoder. In this case, the recommended workflow is to use PCM with professional metadata as described above.

DATA RATE EFFICIENCY FOR NGA DELIVERY

The data rate efficiency of AC-4 extends to delivery of Dolby Atmos and multi-language content. The table below shows the recommended minimum data rates for bitstreams carrying standard 5.1 channel audio as a CM with AD, Dolby Atmos audio, bitstreams carrying multi-language content using separate Dialog(s) and M&E components, and the combination of both. Note that the recommended column has rounded up some data rates in an attempt to minimize changes required when the Channel Config switches.

Channel Config	Recommend Data Rate
5.1 CM + AD	192 kbps (144 + 48)
5.1 M&E+D	192 kbps (144 + 48)
5.1 M&E+D+D	240 kbps (144 + 48 + 48)
5.1 M&E+D+D+AD	288 kbps (144 + 48 + 48 + 48)
5.1.2 CM	224 kbps
5.1.2 M&E+D+D	320 kbps (224 + 48 + 48)
5.1.2 M&E+D+D+AD	368 kbps (224 + 48 + 48 + 48)
5.1.4 CM	288 kbps
5.1.4 M&E+D	336 kbps (288 + 48)
5.1.4 M&E+D+D+AD	432 kbps (288 + 48 + 48 + 48)
7.1.4 CM	320 kbps
7.1.4 M&E+D	368 kbps (320 + 48)
7.1.4 M&E+D+D+AD	464 kbps (320 + 48 + 48 + 48)

AC-4 encoding products for NGA

Dolby develops and delivers different versions of the AC-4 encoder SDK (software development kit) to encoder manufacturers such as those referenced above. All AC-4 encoders, regardless of the version number, support existing ATSC v1.0 use cases (encoding of up to 5.1 channels of PCM into AC-4 5.1 bitstream, RTLL, and Dialog Enhancement analysis). However, features such as Dolby Atmos and efficient multi-language delivery through compositional audio streams (e.g. M&E + D + D) will be enabled by newer versions of the AC-4 encoder SDK.

MVPD contribution

In many cases, broadcasters have direct feeds to MVPDs (multichannel video programming distributors) where baseband video and PCM audio are provided by the broadcaster and encoded by the MVPD for compatible carriage in their systems. MVPDs are also receiving ATSC 1.0 broadcast streams off the air for redistribution on their service. This typically does not require transcoding of the content. However, for ATSC 3.0 signals received off-air by MVPDs, there are critical differences that will, in the earliest days of ATSC 3.0, make transcoding a necessity.

At the transport layer, MVPDs commonly use MPEG-2 transport streams while ATSC 3.0 in North America uses ROUTE (Real-time Object delivery over Unidirectional Transport) content packaged in MPEG DASH (Dynamic Adaptive Streaming over HTTP) segments as the transport layer. At the elementary stream layer, ATSC 3.0 uses HEVC video encoding and AC-4 audio encoding. MVPDs may thus be required to re-package the incoming ATSC 3.0 feed into MPEG-2 transport streams for STBs that support HEVC video and AC-4 audio. MVPDs with STBs that don't support HEVC and AC-4 will also need to transcode the video and audio into formats their STBs can support — such as H.264 or MPEG-2 for video, and AC-3 or E-AC-3 for audio. For MVPD set top boxes that support only MPEG-2 video and AC-3 audio, it is recommended to use the ATSC 1.0 audio, if available. Preferably this would be in the form of PCM audio which is then encoded to AC-3 by the MVPD, or by using the AC-3 encoded version in a transport stream supplied by the broadcaster. This removes the need to transcode the low bit rate AC-4 audio to AC-3, which could degrade quality.

When AC-4 delivers audio in Compositional Audio Format (e.g. multiple presentations via M&E+D1+D2+AD), each of the presentations intended to be carried through an existing MVPD needs to be decoded from the AC-4 multi-presentation bitstream and rendered to a CM, then re-encoded as separate CMs using the codec supported by the MVPD.

It is worth noting that due to legacy infrastructure constraints, often only a primary and secondary audio program are supported for each video program, and the secondary program may be limited to stereo or mono. The AC-4 decoder can be configured to output a CM of the desired presentation (e.g. M&E+D1) in the desired channel configuration (5.1, 2.0, or 1.0), but there might be a forced choice between carrying a main program plus an alternate language, or a main program plus video description, if only two can be supported.

At the time of this publication, only one ATSC 3.0-capable IRD (Integrated Receiver/Decoder) with AC-4 decoding capability is commercially available (DS Broadcast BGD4100) which can provide PCM audio to an external AC-3 or E-AC-3 encoder. The AC-4 to AC-3 or E-AC-3 transcode functionality can also be accomplished using the AC-4 pro-decoder and AC-3/E-AC-3 encoder SDKs, which are available.

Sample workflow

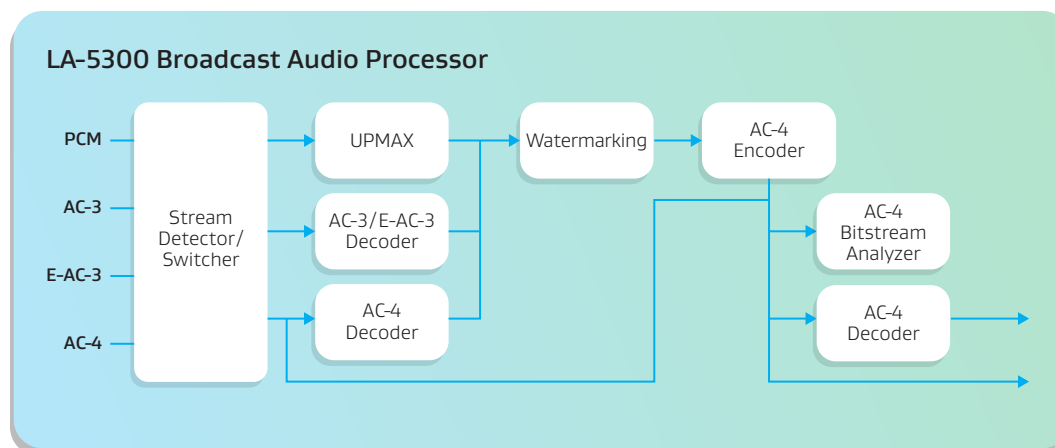
Below is an example of an audio workflow for a local ATSC 3.0 television station. A Linear Acoustic LA-5300 is used for common audio processing functions including upmixing local content, inserting watermarks such as those from Nielsen and Verance, and AC-4 encoding. The AC-4 signal is then applied to a Harmonic Electra X2 for video encoding and AC-4 MPEG-DASH packaging.

In this example, the HD-SDI signal represents a suggested first step approach to carrying advanced next generation audio, such as 5.1.2 M&E+D1+D2+AD, using only one extra pair of audio channels. To accomplish this, the HD-SDI signal uses channels 1–6 to carry a legacy 5.1 or 2.0 channel CM mix in English, channels 7–8 for carrying SAP and/or VDS, and channels 15–16 to carry a pre-encoded AC-4 signal that is the 5.1.2 M&E+D1+D2+AD NGA version.



The LA-5300 transcodes input signals to PCM when necessary, inserts watermarks if desired, and encodes the result to AC-4 for output. When AC-4 is present on input channels 15-16, the LA-5300 either decodes the AC-4 audio stream to add watermarking, or it passes the pre-encoded AC-4 to the output, replacing the local AC-4 encoded signal.

This workflow allows a broadcaster to continue working with PCM audio to feed their emission encoder while automatically passing through advanced AC-4 encoded signals when available.



Additional information

AC-4 decoding in ATSC 3.0 TVs and STBs

ATSC 3.0 televisions and STBs equipped with AC-4 decoding use the Dolby Multi-Stream Decoder MS12. MS12 is a single-package solution for decoding AC-4 audio streams and other audio formats. There are several MS12 versions that support AC-4 decoding.

All AC-4 decoder versions within the MS12 family can

- decode AC-4 audio streams with channel formats ranging from mono to Dolby Atmos, and downmix to stereo, when needed, for playback through the built-in speakers in the television set;
- and decode compositional audio streams (e.g. M&E+D+D), playing back the presentation selected by the user.

For example, all MS12-based AC-4 decoders can decode an AC-4 audio stream containing M&E + English Commentary + Spanish Commentary, and produce a final mix combining the M&E element with either language as chosen by the end-user through the television's user interface.

However, in order to deliver the full value of the ATSC 3.0 platform, an AC-4-enabled TV or STB must offer these two features, not present in all MS12 implementations.

DOLBY ATMOS OUTPUT

This feature enables output of Dolby Atmos audio from the TV or STB when the source AC-4 audio stream contains a Dolby Atmos mix. This is necessary in order to offer a full Dolby Atmos experience through a connected Dolby Atmos AVR or Soundbar when the viewers connect their ATSC 3.0 receiver to this type of audio playback device.

NON-LANGUAGE-BASED PRESENTATION SELECTION

All AC-4 decoders within the MS12 family can select the necessary elements from a compositional AC-4 audio stream (e.g. M&E+D+D) in order to build a full mix. However, only MS12 implementations v2.3 and later support name tags for presentations that are not languages (e.g. English/Spanish) and/or Visually Impaired presentations. Other presentations using the same language but different content — for example, a sporting event with separate user-selectable commentaries for home and away teams — are not supported.

MS12 v2.3 and newer meet the minimum requirement to provide the Next Generation Audio experience that ATSC 3.0 offers.

- MS12 v2.3 and newer can output a Dolby Atmos mix from the TV or STB output to a Dolby Atmos AVR or Soundbar.
- MS12 v2.3 and newer supports a wider selection of name tags for the different audio presentations.

Emergency audio

Emergency audio includes two primary functions: Emergency Alert and Emergency Information.

Emergency Alert is the well-known system where normal program audio is replaced with emergency alert tones and messages. This is most often accomplished today prior to audio encoding and is kept simple and direct for highest reliability of this critical function. ATSC 3.0 does not vary from this model — current functionality continues: normal program audio is replaced with emergency alert audio prior to the AC-4 encoder. If a broadcaster is using pre-encoded AC-4 from network distribution or a playout server, it is also feasible to accomplish this by replacing the pre-encoded AC-4 signal with locally AC-4 encoded emergency alert audio. In either case, emergency audio is delivered to all decoders, regardless of any preference settings consumers may have chosen.

Emergency Information is new to ATSC 3.0 and essentially enables important messages that are displayed in the video to also be made available aurally as Associated Audio for consumers who might be visually impaired and have set the preferences in their televisions to reproduce associated audio. An example is text-to-speech conversion of an information crawl; e.g. traffic conditions. The text-to-speech audio is encoded as an associated audio substream of type Emergency (E) and corresponding presentations are added. Receivers with preferences set to reproduce associated audio will attenuate the main program audio to prioritize output of the associated Emergency Information audio. If the receiver preferences are not set this way, or the visually impaired user disables the Emergency Information temporarily, output of the main program audio will continue.

Ad insertion

ATSC 3.0 defines a system for targeted advertisement that is based on a client-side ad insertion model (CSAI). ATSC 3.0 broadcasters can provide a browser-based application to manage the insertion of the targeted ads, as well as the consumption data gathering and reporting. The broadcaster application can access a dedicated API defined by ATSC 3.0 for the purpose of querying receiver information. This API allows the application to:

- Evaluate the receiver information
- Receive notifications from broadcast signaling or client status
- Request the receiver to perform actions

ATSC 3.0 defines “ad-avails” — markers delivered with the broadcast content used to indicate an advertisement placement opportunity to the client. Ad-avails enable substitution of broadcast-delivered non-targeted advertisements by placing Periods with XLinks in the DASH MPD delivered within the ATSC 3.0 broadcast. XLinks can be used to reference external targeted advertisement content. The receiver’s DASH Player will identify the location of the targeted advert and handle the substitution of the non-targeted broadcast advert. Adverts for ATSC 3.0 have to be DASH/ISOBMFF formatted; the DASH/ISOBMFF segments (XLink-resolved URLs) that make up the targeted advert content can be downloaded prior to playback. As ATSC 3.0 in North America requires the use of AC-4 audio, targeted adverts must be encoded in AC-4. To allow very fast switching between broadcast and OTT-delivered (Over The Top) targeted ads, the targeted advert audio content should:

- Use the bitstream configuration that best matches the broadcast content
- Ensure that the AC-4 bitstream is encoded at a frame rate that matches the video
- Ensure that the first AC-4 frame of the targeted advert audio is an AC-4 I-frame
- Ensure that the first AC-4 frame of the broadcast audio after the Period has passed also contains an AC-4 I-frame
- Ensure that ISOBMFF segmentation structure of audio and video are aligned

These recommendations will enable the receiver to switch between targeted ad audio and broadcast audio without needing to re-initialize the AC-4 decoder or other audio processing elements.

The ATSC 3.0 specification also allows the use of watermarking technology instead of DASH XLinks to indicate the availability of ads to an ATSC 3.0 client.

These documents detail ATSC 3.0 ad insertion specifications:

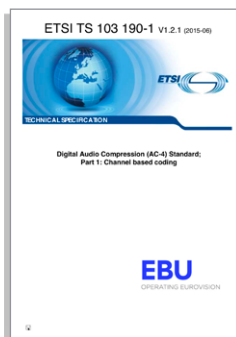
- [ATSC A/344](#) (description of communication between entities in a TV that manage ad insertion)
- [ATSC A/337](#) (description of the use of events for notifications)
- [ATSC A/334](#) (specifies the VP1 audio watermark for systems conforming to the ATSC 3.0 family of specifications; emission by broadcasters is optional)
- [ATSC A/335](#) (describes video watermark technology to robustly embed ancillary data in the transmitted pixels of a video signal; emission by broadcasters is optional)
- [DASH-IF IOP for ATSC 3.0](#) (DASH signaling)
- [DASH-IF IOP](#) (more detailed ad-insertion use case description, DASH signaling, event handling, etc.)

Getting started

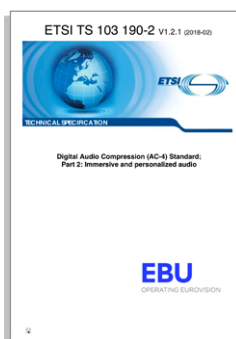
ATSC 3.0 is on air today with AC-4, delivering Next Generation Audio experiences and enabling existing services to be transmitted with better quality and efficiency. This handbook has provided an overview of the core technical considerations and implementation paths that need to be considered in designing an ATSC 3.0 transmission plant. The industry, with ongoing assistance from Dolby, continues to develop upstream parts of the ecosystem required to provide more NGA features.

Resources

ETSI

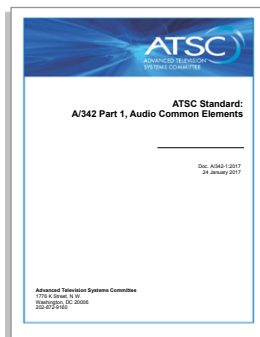


TS 103 190-1 V1.2.1

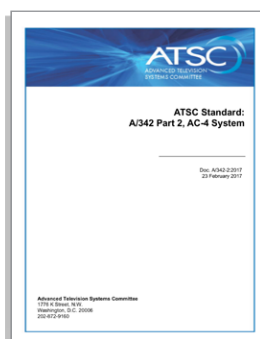


TS 103 190-2 V1.2.1

ATSC

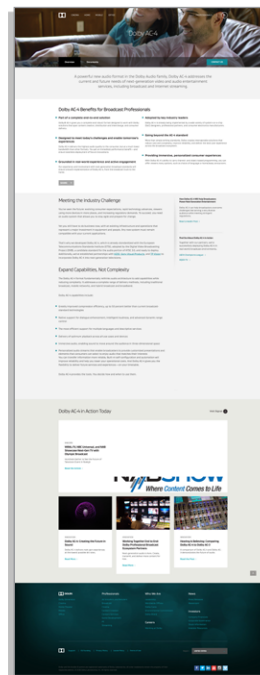


A/342 Part 1



A/342 Part 2

DOLBY



AC-4

Abbreviations

AC-3	Audio Codec 3
AC-4	Audio Codec 4
ATSC	Advanced Television Systems Committee
AD	Audio Description
CSAI	Client-Side Ad Insertion
CM	Complete Main
DASH	Dynamic Adaptive Streaming over HTTP
E-AC-3	Enhanced Audio Codec 3
ED2	Enhanced Dolby E
IRD	Integrated Receiver/Decoder
LFE	Low-Frequency Effects
M&E + D	Music and Effects Plus Dialog
MADI	Multichannel Audio Digital Interface
MPEG	Moving Picture Experts Group
MVPDs	Multichannel Video Programming Distributors
NGA	Next Generation Audio
OB	Outside Broadcast
PMD	Professional Metadata
PCM	Pulse Code Modulation
ROUTE	Real-time Object Delivery Over Unidirectional Transport
RTLL	Real-Time Loudness Leveler
SoC	System on Chip
STB	Set Top Box
SDI	Serial Digital Interface
SDK	Software Development Kit



For more information, please contact us at **DolbyAC-4Inquiry@dolby.com**