

B DVD-Audio Format

Like DVD-Video, the DVD-Audio format is built on the foundation of DVD-ROM. A DVD-Audio disc is actually a DVD-ROM disc that includes a zone of material supported by the DVD-Audio specification (and may optionally contain a zone of DVD-Video material). These zones are essentially directories or folders in the Universal Disk Format (UDF) file system.

The DVD-Audio specification incorporates many of the navigational and architectural features of DVD-Video, but also includes several interactive and playback features. It also supports higher audio fidelity and more flexible multi-channel playback options. While the two formats are closely related, DVD-Audio is a truly distinct product tailored for audio-centric applications.

For those intending to create DVD-Audio albums with DVD-Audio Creator, this appendix outlines important information that can help you get the most out of the format. It also clarifies the functional and technical similarities and differences between DVD-Audio and its DVD-Video cousin. It includes the following topics:

- “DVD-Audio Players” on page 192
- “DVD-Audio Content” on page 195
- “The Structure of DVD-Audio” on page 208
- “Using Video in DVD-Audio” on page 212
- “DVD-Audio Title Example” on page 215

DVD-Audio Players

The specifications for DVD-Audio were developed by Working Group 4 (WG-4) of the DVD Forum, a consortium representing companies involved in consumer electronics, computing, and entertainment. To ensure DVD-Audio's flexibility across a wide range of applications, WG-4 decided not to define one single set of features that would be mandated for all DVD-Audio players. Instead, a DVD-Audio disc can contain a variety of different content types, each of which will play on one or more of several different player types designed for different playback settings.

While this approach sounds complicated, much of the distinction between the different envisioned playback settings boils down to the question of whether or not video display is available. With an audio format, it is crucial to support playback without a video display, particularly with portable *Walkman* devices, *boomboxes*, and cars. Even for home listening, requiring video support in all players would increase manufacturing costs, making it harder to reach the less-expensive end of the audio player market.


Though the need for audio-only playback was clear, DVD-Audio's developers also needed to support all the visual extras that consumers would expect from a next-generation entertainment format: graphics, text and motion video. To cover all bases, WG-4 envisioned that five roughly-defined player types might possibly be used to play content from a disc in the DVD-Audio format. These include:

Two audio-only players (AOPs):

- A *simple* audio-only player has no video output, and navigates the disc linearly using a list of tracks that is analogous to the table of contents (TOC) on a CD.
- A *smart* audio-only player has no video output, but gives the user more choice of how to navigate the material on the disc, and may include an LCD display to show song titles and other text information such as lyrics.


Video-capable players (VCAPs):

- An Audio-with-video player supports navigational choice and includes video outputs to support multimedia, including visual menus, album artwork, slideshow graphics, and motion video (such as music videos).
- A universal player offers the same audio and multimedia support as the Audio-with-Video DVD-Audio player, but also plays DVD-Video discs.
- A DVD-Video player can be used to play optional video content that has been included on a DVD-Audio disc (though the DVD-Audio content on a DVD-Audio disc will not play on a DVD-Video player).

**Figure 2:** *DVD-Audio discs and player types*


While the different player types shown in Figure 2 may be inferred from the DVD-Audio specification, it is considered unlikely that all of them will actually reach market. In particular, it is unclear what demand there might be for the Audio-with-Video player that is not a Universal player. Similarly, many believe that there will be little justification for a distinct video-only player once the Universal player is widely available at a competitive price. Eventually, new consumer DVD players may well be sold in just two main categories: Audio-only and universal.

One way to easily visualize the difference between the two main categories of DVD-Audio players (Audio-only and Universal) is to compare the typical remote control implied for each type in the specification. The Universal player's remote control looks very much like that of a DVD-Video player:



Remote control for a Universal DVD player

In contrast, the minimal set of functions required to control the Audio-only player makes for a very simple remote:



Remote control for an Audio-only DVD-Audio player

DVD-Audio Content

While the availability of different player types allows DVD-Audio discs to be played in a variety of settings, it also means that it is possible for a DVD-Audio disc to include some types of content that will not be available for playback in some situations. By choosing which types of content are put on a given DVD-Audio title, it is the producer who decides what the user's experience will be in various playback environments.

Linear audio-only discs The base-case DVD-Audio disc is one that contains no graphics, text, or motion video information. This *Pure Audio* disc supports higher fidelity, greater capacity, and more channels than a CD, but it functions in much the same way, with similar track-based navigation. Pure Audio discs play on both Audio-only and Universal DVD-Audio players, but they do not play on DVD-Video players.

Navigable audio-only discs Somewhat more sophisticated DVD-Audio titles can be created by using the format's logical structure to offer listeners different navigational options, such as multiple paths (playlists or *groups*) through the audio material. These discs play on simple Audio-only players, but their more sophisticated features come to life when played on smart Audio-only players, or on video-enabled players such as Universal players. They do not play on DVD-Video players.

Multimedia discs More complex DVD-Audio titles take greater advantage of the format's multimedia capabilities, using visual menus for navigation, slideshows with audio, and motion video. These titles play audio on Audio-only players, but without all the extras that are viewable on video-capable players such as Universal players. If such a disc includes motion video content, that content (only) is playable on DVD-Video players.

The following table provides an overview of some common combinations of disc content and player types:

DVD-Audio Player Type	DVD-Audio disc without video	DVD-Audio disc with video	DVD-Video disc
Audio-only player	<ul style="list-style-type: none"> • Audio • Text information 	Audio elements of DVD-Audio and DVD-Video titles	Not playable
Universal player	<ul style="list-style-type: none"> • Audio • Still pictures • Text information 	<ul style="list-style-type: none"> • Visual menus • All audio, text, graphic and video content 	All DVD-Video content supported
DVD-Video player	Not playable	Video content only	All DVD-Video content supported

Audio Formats

The core content type of the DVD-Audio format is, obviously, audio.



While the DVD-Video specification gives equal weight to both linear PCM (pulse code modulation) and compressed audio formats such as Dolby Digital (AC-3), the DVD-Audio specification puts the priority on delivering the highest possible audio fidelity, and thus focuses primarily on PCM.

PCM data can be stored in either of two forms: linear (LPCM) or *packed* using Meridian Lossless Packing (MLP). All players are required to support both LPCM and MLP audio in either stereo or multi-channel configurations. Players with video outputs must also support Dolby Digital. Optional audio formats such as DTS, or others not yet defined, may be supported at the manufacturer's discretion.

DVD-Audio content can offer up to two audio streams in any given program. The first stream is required to be PCM (stereo or multi-channel), the optional second stream can be an optional audio type.

Audio Resolution

With the PCM approach (LPCM or MLP), DVD-Audio offers content developers a wide array of choices regarding resolution. In PCM, the resolution of a given audio stream depends on two factors: the word length and the sample rate. The sample rate is the number of times per second that the amplitude of the source analog waveform is sampled for conversion into a digitally stored value. The word length is the number of bits available to store the value of each sample.



Audio resolution (fidelity) is improved with higher sample rate and word length

DVD-Audio’s PCM support covers two families of sample rates. One family is based on the 44.1 kHz rate of the audio CD, and also includes the multiples 88.2 and 176.4 kHz. The other family starts with the 48 kHz rate commonly used with DAT and video, and adds 96 and 192 kHz.

DVD-Audio supports several word lengths, including 16- (used in audio CDs), 20- and 24-bit. A longer word length translates into finer resolution and a wider dynamic range.

Multi-Channel Sound

Of all of DVD-Audio’s features, one of the most appealing to consumers is the format’s ability to deliver multi-channel sound. In each sample-rate family, the highest sample rate (176.4 or 192.0 kHz) supports only mono or two-channel playback. The remaining rates (44.1, 48, 88.2, and 96 kHz) support mono through six channels.

Channels supported for each sample rate with DVD-Audio

	44.1 kHz Family			48 kHz Family		
Sample Rate (kHz)	44.1	88.2	176.4	48	96	192
Maximum Channels	6	6	2	6	6	2

DVD-Audio’s maximum data rate allocable to audio—9.6 Megabits per second—places limits on the total data bandwidth available to the channels. Six channels of 20-bit audio at either 88.2 or 96 kHz, for example, would substantially exceed the format’s data rate. The specification offers two complementary strategies for dealing with this limitation. One is MLP (see “Playing Time and Meridian Lossless Packing” on page 201), the other is mixed resolutions.

Mixed resolutions allow the producer to prioritize the allocation of bits among the channels in a given audio track. Each channel is assigned to one of two channel groups, with the resolution of Group 1 always equal to or greater than that of Group 2. Within a given track, the sample rates used for all channels in both groups must be from the same family.

Supported sample rates and word lengths for DVD-Audio's two channel groups

	If Channel Group 1 is:	... then Group Channel 2 may be:
Sample rate	48 kHz	48 kHz
	96 kHz	96 or 48 kHz
	192 kHz	
	44.1 kHz	44.1 kHz
	88.2 kHz	88.2 or 44.1 kHz
	176.4 kHz	
Word length	16 bits	16 bits
	20 bits	20 or 16 bits
	24 bits	24, 20 or 16 bits

Note: With a sample rate of 192 kHz or 176.4 kHz, the number of channels must be two or less (two-channel stereo or mono).

Assignment of channels to groups, and selection of attributes within each group, can be changed on a track-by-track basis. To regularize the use of mixed resolutions among channels, most, but not all, of the reasonably foreseeable possibilities for varying the resolution of up to six channels are defined in the format's 21 supported channel configurations.

Table 2 shows the supported channel assignments. Assignment 8 uses three channels, with priority to left and right. Assignment 21 uses six channels, with priority to the four corners (left and right front, left and right surround).]

Table 2: *DVD-Audio channel assignments for two groups*

		ch0	ch1	ch2	ch3	ch4	ch5
	1	C					
Priority to front L and R	2	L	R				
	3	L	R	S			
	4	L	R	Ls	Rs		
	5	L	R	Lfe			
	6	L	R	Lfe	S		
	7	L	R	Lfe	Ls	Rs	
	8	L	R	C			
	9	L	R	C	S		
	10	L	R	C	Ls	Rs	
	11	L	R	C	Lfe		
	12	L	R	C	Lfe	S	
	13	L	R	C	Lfe	Ls	Rs
Priority to front L, R, and C	14	L	R	C	S		
	15	L	R	C	Ls	Rs	
	16	L	R	C	Lfe		
	17	L	R	C	Lfe	S	
	18	L	R	C	Lfe	Ls	Rs
Priority to corners: L, R, Ls, Rs	19	L	R	Ls	Rs	Lfe	
	20	L	R	Ls	Rs	C	
	21	L	R	Ls	Rs	C	Lfe
Channel Group 1						Channel Group 2	

To illustrate the range of channel groupings permitted by the DVD-Audio specification, Figure 3 shows some examples shown with possible attribute assignments.




Figure 3: Supported channel assignments for two DVD-Audio's two channel groups

Playing Time and Meridian Lossless Packing

Because DVD-Audio supports so many different variations of sample rates, word lengths, and channel configurations, the capacity in minutes of the format varies widely. The ability to vary these attributes on a track-by-track basis adds to the complexity of stating the maximum playing time. But it is clear that the use of more channels, more samples, or more bits per sample all result in shorter playback capacity.

Table 3: Playing times for DVD-Audio variations

Quantization (bits)	Sample Rate (kHz)	Channels	Data Rate (mbps)	Time (minutes)
16	44.1	2	1.41	422
16	48	2	1.54	388
20	96	4	7.68	78
24	192	2	9.22	65

Note: Single-layer, single-sided disc with no lossless compression; assumes 5% overhead for navigation and formatting data.

As shown in the Table 3, at the same resolution as CD-Audio (stereo 16-bit/44.1 kHz), DVD-Audio is capable of far greater playing time than CD-Audio's 74 minutes. But even with DVD-Audio's greater data rate and data capacity, full support for high-fidelity surround sound (six channels of 24-bit/96 kHz audio, for instance) would exceed the format's maximum bit rate and quickly use up the available storage.

To address this issue, WG-4 adopted the Meridian Lossless Packing (MLP) compression system developed by Meridian Audio. Because it is a completely lossless system, a decoded MLP data stream is bit-for-bit identical to the pre-encoded PCM source stream.

MLP permits substantial reduction in the bandwidth required to store high-quality multi-channel PCM audio. The efficiency of the algorithm varies with the program content, particularly for material at sample rates of 44.1 or 48 kHz. For higher sample rates (88.2, 96, 176.4 and 192 kHz), however, MLP has been demonstrated to reliably achieve a 45% reduction in bandwidth requirements. This allows a DVD-5 disc to be used for presentation of up to 74 minutes of audio using six channels of 24-bit/96 kHz audio. DVD-Audio discs are not required to use MLP, but decoding capability for MLP is mandated for all DVD-Audio players.

SMART Content and Downmixing

One of the problems confronted by WG-4 in defining DVD-Audio stems from the fact that, aside from home theater systems, most playback situations (stereo receivers, headphones, and boom-boxes, for example) are currently configured for stereo. To ensure that a DVD-Audio disc will always play, regardless of the available monitoring configuration, WG-4 had to define how a player should present multi-channel programs in situations where only stereo playback is available. The difficulty is that a program mixed for surround presentation will not sound correctly balanced if the player simply ignores all channels other than front left and right.

DVD-Audio allows both stereo and multi-channel mixes of the same music to be delivered on one disc. One solution would have been to require that every disc that includes a multi-channel program must also include a two-channel version of the same program. But that would have imposed significant limitations on the playing time of discs with multi-channel programs.

Instead, WG-4 mandated that players support an approach championed by Warner Bros. Records called “SMART Content” (System Managed Audio Resource Technique). SMART allows the producer to determine in advance a set of downmix parameters (coefficients) defining the relative level, panning and phase that will be applied to each channel of a multi-channel mix if it is combined into stereo. A SMART downmix is only played if a discrete 2-channel mix of a given program has not been included on the disc.

SMART coefficients are specified as 8-bit numbers in a six-by-two matrix, along with a phase bit for each channel. A set of 16 tables of these coefficients can be defined for each Audio Object Set. The table for each individual track of a title can be selected from among these 16. This allows producers a workable means to control the sound of their music in cases where their multi-channel mix is downmixed by the player.

The Audio Object (AOB)

Like the other presentation data (pictures, video, and text) that may be included in a DVD-Audio album, the audio data is stored in files that are referenced in the disc’s logical data files (which determine the conditions under which the audio data files are actually played).

The basic unit of audio presentation data is the Audio Object (AOB), a file containing the audio data for one or more selections (songs or compositions). In each AOB, all audio attributes (resolution, channel configuration, etc.) remain consistent throughout.

An AOB can also contain Real-Time Information (RTI) that can be used to display Real-Time Text if that feature is supported by the player (see “Text Information” on page 206).

Value-Added Content

While audio playback is at the heart of DVD-Audio, the format is also intended to offer a workable platform for the kinds of value-added multimedia features that have proven difficult to deliver to the mass market on formats such as CD-ROM and Enhanced CD. These features include graphical accompaniment for the music, text information about the music and the artist, and motion video such as music videos or interviews.

Still Pictures

The DVD-Audio specification supports the display of still pictures during audio playback, and offers the producer a wide range of options as to when and how the pictures are shown.

Still pictures in DVD-Audio are MPEG-2 encoded images. There are two basic categories of pictures:

Slide A picture with no buttons and comprised of a single image layer.

Menu A picture with up to 36 buttons and compromised of up to three layers, which include:

- A 24-bit background layer containing the visible background image
- A subpicture Overlay (SPU) layer of up to four colors for shapes (including text) that are displayed over the background image
- A button hotspots layer for highlight information (HLI), which defines the rectangular boundaries of each of the menu's buttons

A still picture that is displayed during playback of an audio track is referred to as an ASV (Audio Still Video). While the data size of an individual ASV may vary greatly, typical ASVs are approximately 100 KB.

ASVs are grouped into sets of graphics referred to as ASVUs (Audio Still Video Units), each of which can contain up to 99 ASVs (individual images). The audio selections over which the ASVs from a given ASVU are intended to display is referred to as the ASVU range. An ASVU range can be either a single track or a set of tracks.

The total size of an ASVU may not exceed 2 Megabytes, which is the size of the buffer that player manufacturers are required to provide to allow ASVUs to be preloaded into player memory. This preloading process ensures that players will not interrupt audio playback in order to read graphical data from the disc. For instance, a listener can browse a series of still graphics without affecting the currently playing song.

From an authoring point of view, it is crucial to realize that audio output is muted during preloading of ASVUs. Depending on player design, the duration of the muting between each ASVU range is at least a second, possibly more. This limitation means that the use of graphics must be carefully planned in situations such as live albums, where audio may continue for long periods without the opportunity to load a new ASVU.

Modes, Effects, and Highlights

For each ASVU, the producer defines a playback mode that determines the order and duration of graphical display:

Slideshow ASVU

The display duration for slideshow ASVs is predetermined during authoring and is not changeable by the user. There are three types of slideshows:

Sequential ASVs are displayed in a predetermined order.

Random ASVs are displayed in an order generated at random by the player during playback.

Shuffle Same as random, but no ASV is repeated until every ASV has been shown.

Browsable ASVU

The display duration for browsable ASVs can be affected by user input from the remote control. There are four types of Browsable ASVUs:

Timed sequential ASVs are displayed in a predetermined order. The display duration of each ASV is either fixed (predetermined) or randomly generated by the player from within a predetermined range. The user can override the duration and browse the ASVs by using the PAGE (+/-) keys on the remote control.

Timed random ASVs are displayed in an order generated at random by the player during playback. The display duration of each ASV is either fixed (predetermined) or randomly generated by the player from within a predetermined range. The user can override the duration and browse the ASVs by using the PAGE (+/-) keys on the remote control.

Timed shuffle Same as timed random, but no ASV is repeated until every ASV has been shown.

Infinite sequential ASVs are displayed in a predetermined order. The display duration of all ASVs is predetermined as infinite (displayed until there is user input from the remote control). The user can override the duration and browse the ASVs by using the PAGE (+/-) keys on the remote control.

The DVD-Audio specification supports a variety of start and end effects (wipes and dissolves) as transitions between ASVs.

The specification also allows the highlights within a given ASV to change over time. This is useful for content such as song lists, lyrics or libretto. Links allowing listeners to navigate to different parts of a song by clicking on lyrics are an appealing feature of DVD-Audio titles (see “Linked Lyrics” on page 57).

Text Information

The DVD-Audio specification provides for storage of text information on disc, though inclusion of text is not mandatory.

In the specification, text support is provided for multiple “language units,” with up to eight languages available on any given disc. Two character types are supported. For languages such as English that use single-byte characters (ISO8859-1), text is structured for presentation as *pages* of four lines of up to 30 characters each. Japanese and other languages using double-byte characters (JIS Kanji) are organized into pages of two lines of up to 15 characters each.

The manner in which any given player uses this text information (if at all) is left to the discretion of the player manufacturer. A player with video output, for example, can display the text on-screen. But the size of the pages defined by the specification suggests that the primary application is envisioned as an LCD screen on *smart* Audio-only players, Audio-with-Video players, and Universal players.

The specification defines two distinct types of text:

AudioText Data Intended for display of static information that is not synchronized to the audio program. If Audio Text Data is used, it must include the album title, group name, and track name. It can also include optional information such as the names of the artist and composer. Up to 64 KB of data can be included for each language unit.

RealTimeText Stored on disc as part of the audio stream, allowing synchronization with audio playback. Ideal for lyrics and libretto, Real Time Text can also be used for context-dependent commentary such as running liner notes.

The extent to which Real Time Text or Audio Text Data is supported by a given player is up to the hardware manufacturer. As of this writing, no players on the market support the format's text display capabilities.


Video on DVD-Audio

DVD-Audio's support for motion video is similar to that of DVD-Video, using Video Objects (VOBs) comprised of MPEG-2 video plus audio and optional subtitles. However, some advanced features that are supported in DVD-Video VOBs—such as parental control and complex interactivity—are not supported in all DVD-Audio settings. To fully understand these limitations, it is necessary to understand the overall organization of the DVD-Audio format (see “The Structure of DVD-Audio” on page 208).

The Structure of DVD-Audio

A DVD-Audio disc is a DVD-ROM disc that includes a special zone (the Audio_TS folder or directory) for DVD-Audio material. A DVD-Audio disc can also include a Video_TS folder for an optional DVD-Video zone, as well as a DVD-Others zone for applications and data that can be accessed from a computer-hosted DVD-ROM drive.

The overall volume structure of the DVD-Audio disc requires that the DVD-Audio area appear first, followed by the DVD-Video area and any other non-audio and video files and directories.



Volume structure of a DVD-Audio disc

The final playback experience of a given DVD-Audio title involves two interrelated elements of DVD playback: the presentation data (audio, still pictures, text and motion video) and the logical data that determines the order and conditions of playback. Because DVD-Audio discs are intended for playback on a number of different player types, the way in which the information on disc is utilized and experienced also depends on the type of player.

Multiple Content Managers

In the world of DVD, the first place a player goes to find out what content is present on a given disc is a directory (table of contents) referred to as a *Manager*.

For playback in a DVD-Audio player, a disc is required to include in the Audio zone several different content directories, each used to accommodate a different playback situation (player). A DVD-Audio disc that includes an optional DVD-Video zone also includes a Manager in the Video zone (a Video Manager or VMG), but this directory is only used if the disc is played in a DVD-Video player.

The most basic form of Manager in DVD-Audio is the SAMG (Simple Audio Manager). Containing a list of up to 314 tracks, it serves the same function as a table of contents on a CD-Audio disc. Simple Audio-only players look at SAMG to find the information they need for linear, track-based navigation of the disc's audio content.

Smart Audio-only players, Audio-with-Video players, and Universal players all support more sophisticated navigation than do simple Audio-only players. The directory information these players need is found in the Audio Manager Information (AMGI) section of the AMG (Audio Manager).

The AMGI contains several subsections, two of which are tables of *Search Pointers* telling the player the location on disc of all content that is playable by that type of player. Smart Audio-only players use the AMG Search Pointer table designated as AOTT_SRPT, while video-capable players use the table ATT_SRPT.

In addition to these Search Pointer tables, AMG contains the logical data and presentation data (audio, still pictures, overlays, video) for the Audio Manager Menu (AMGM), which is the visual menu that is the destination of the TOP MENU key on the player remote control. The data making up the AMGM is referred to as the *Audio Manager Menu Domain* (AMGM_DOM).

Logical Structure: Albums, Groups, and Tracks


Except for AMGM, all DVD-Audio content makes up the *Title Group Domain* (TT_GR_DOM). DVD-Audio players that read AMG (all except simple Audio-only players) are able to take advantage of DVD-Audio's capacity to organize this content hierarchically rather than linearly. This logical hierarchy is based on five different levels: album, group, title, track and index.

Each side of a DVD-Audio disc contains one album. Each *album* can contain up to nine groups, each of which is essentially a playlist specifying the playback order of a number of titles. While any title can contain up to 99 tracks, there may be no more than 99 tracks total within a single group.

As on a CD, a *track* can be thought of as an identifiable chunk of program (such as a song). A track's presentation data can be either audio or video. During playback, audio track attributes such as channel configuration, sample rate, and word length can be changed on track boundaries (players mute during such attribute changes). Audio track playback can be accompanied by still pictures (ASVs in an ASVU).

An *index* is a reference point to a portion of an audio track (a cell). There can be up to 99 indexes within a single track.

While this hierarchy has many levels, in practice the user is never aware of the existence of titles. On a smart Audio-only player, for example, a user would normally be able to play any track using the remote by simply entering a group number and then the number of a track within that group.



DVD-Audio album with three groups

Because a *group* is simply a playlist, a given track can be referenced by more than one group. On an album with twenty audio tracks, for example, one group might be a sequence of all the songs, another could be a “mellow” playlist of just acoustic numbers, and a third might be a “party” playlist of just dance tracks. Groups thus allow producers to create up to nine different listening experiences drawn from one underlying set of material.

Navigation and Program Chains

DVD-Audio discs are required to include several different content directories to facilitate navigation. The various players that use these directories are characterized as follows:

SAMG players Simple Audio-only players that read SAMG and ignore AMG.

AOTT players Smart Audio-only players that read the AOTT_SRPT table in the AMG.

ATT players Video-capable players (Audio-with-video and Universal players) that read the ATT_SRPT table in the AMG.

Navigationally, the difference between ATT players and AOTT players is that with ATT players the content can be navigated with visual menus like those used in DVD-Video. AOTT machines can't use visual menus because they have no video output. In both cases, however, the underlying structure being navigated is the same:

- A set of groups is defined in the AMG's Search Pointer table
- Each group is a list of titles, with a pointer to the location of each corresponding title set
- Each title set contains:
 - Audio Title Set Information (ATSD), logical data defining the manner in which the title's content will play, and including references to the locations of the Title's presentation data (audio, still pictures, and/or video).
 - Audio Object Set (AOBS), a collection of Audio Object (AOB) files, which are the presentation data for the audio content of titles that do not contain motion video content (AOTT Titles).

The ATSI includes a set of Program Chains (PGCs) that are instructions that tell the player what programs (PGs) to play and when to play them. The content unit played back for each program is a track.

The DVD-Audio specification includes a set of navigation commands for use in audio program chains, as well as General and System parameters that are stored in player memory. As with DVD-Video, the commands are broken into various categories (Link, Jump, GoTo, SetSystem, Set, and Compare). However, because groups are intended to be DVD-Audio's primary mechanism for defining a specific path through the available content on the disc, the navigation commands supported in DVD-Audio are a subset of the DVD-Video set. The availability of nine groups allows navigation that is far more varied than that of CD-Audio, though the possibilities are more limited than with DVD-Video.

Among the most important navigational constraints in DVD-Audio are those relating to movement within a group. A user may enter a group at any point by selecting a specific track. But once playback has started within a given group, that group will continue playing in order through to the end, unless the user exits playback by using the TOP MENU or GROUP button on the remote control. The specification makes no provision for linking directly from a given group to content outside that group.

Using Video in DVD-Audio

DVD-Audio groups are comprised of audio titles that include one or more tracks. If the title is comprised of audio tracks, the title's presentation data is drawn from an AOB and (optionally) an ASVU. In the specification, this kind of audio title is referred to as an *AOTT*.

Groups can also contain audio titles whose presentation data is a Video Object (VOB) stored in a Video Title Set (VTS) in the disc's DVD-Video zone. This kind of audio title is referred to as either an AVTT or an AVTT/AOTT, depending on whether the video soundtrack is enabled for playback on audio-only players (see "Video Sound in DVD-Audio" on page 213).

Whether a VTS is used in an AVTT Title or an AVTT/AOTT Title, the DVD-Audio specification requires that it be playable in a DVD-Video player. Therefore, each VTS on a DVD-Audio is actually referenced in two different managers on the disc: the AMG (Audio Manager) and the VMG (Video Manager).

Video Sound in DVD-Audio

Even though the VOBs used in DVD-Audio are also playable in DVD-Video players, some VTS functionalities supported by the DVD-Video specification are not supported for a VTS used on a DVD-Audio disc. One such area where DVD-Video and DVD-Audio differ is the audio format requirements for video soundtracks.

The DVD-Audio specification requires that every audio track include at least one PCM stream (LPCM or MLP). This rule may or may not apply to the soundtracks of VOBs, depending on the player types for which the video's soundtrack is intended.

During the DVD-Audio production process, the producer decides whether or not the soundtrack of each VTS (if any) on the disc should be playable on audio-only players. This means there are two types of audio titles that draw content from a VTS:

AOTT/AVTT The soundtrack is intended for playback on both audio-only and video-capable players. The VTS must follow the same audio requirements as an AOTT (one mandatory PCM stream plus one optional additional stream).

AVTT The VTS soundtrack is not intended for playback on audio-only players, therefore a PCM stream is not required (either PCM or Dolby Digital may be used).

Requirements for video soundtracks in DVD-Audio

	AOTT/AVTT VTS	AVTT-ONLY VTS
Playable on Audio-only players?	Yes	No
Required Audio Format	Linear PCM	Dolby Digital or LPCM
Optional Audio Format	Dolby Digital	Dolby Digital or LPCM
LPCM Parameters	Max. channels, 48kHz: 8 Max. channels, 96kHz: 2 Bits: 16, 20 or 24 Max. bit rate: 6.1444 Mbps	Same as AOTT/AVTT
Dolby Digital Parameters	Max. channels: 5.1 Max. bit rate: 448 kbps	same as AOTT/AVTT

The DVD-Audio specification requires that if a given type of player is able to play some of the content in a group, it must be able to play all of the content in that group. This means that a group that includes an AVTT Title (a VTS that is not enabled for playback on audio-only players) may not also contain an AOTT Title or an AVTT/AOTT Title (which do play on audio-only players). In effect, videos that are not enabled for (audio) playback on audio-only players must be segregated into their own groups.

Additional Video Rules and Limitations

Another distinction between video on a DVD-Video disc and video on a DVD-Audio disc is that a VTS on a DVD-Audio disc may not include an individual menu (VTSM); the only menu in the Video zone on a DVD-Audio disc is the main Video Manager Menu (VMGM).

Additional differences are found in the capabilities of video PGCs (used to control which VOBs should be played under which conditions). The features available to control VOBs in DVD-Audio are a subset of those available on a DVD-Video disc. The degree to which the video feature set is different depends on the domain in which the video is used.

A Video zone (VIDEO_TS directory) on a DVD-Video disc is made up of two distinct domains. The *menu domain* refers to Video Manager (VMG), which contains a VMG menu (VMGM) with either still image or motion video (VOB) backgrounds. The *title domain* refers to a series of Video Title Sets that normally contain the main video content of a DVD-Video disc. Both the menu and title domains of the Video zone draw their presentation data from VOBs.

If motion video content is present on a DVD-Audio disc, the disc still contains a Video zone with both menu and title domains. However, the menu domain of the Video zone is never seen during playback on any type of DVD-Audio player (including Universal); it is present only to allow the disc's Video Title Sets to be viewed on a DVD-Video player.

In addition to the domains in the Video zone, a DVD-Audio disc has two additional domains that are both in the Audio zone (AUDIO_TS): Audio Manager (AMG) and Title Group. Like VMG menus, the visual menus in the menu domain of the Audio zone (AMGM) may use motion video backgrounds. The attributes of VOBs in the AMG are nearly the same as, but not identical to, those of VOBs on a DVD-Video disc. Because features such as pre- and post-commands are supported, it is possible to create complex interactivity in the AMGM.

The use of video in VTS-based Audio Titles (AVTT and AVTT/AOTT)—which are stored in the Video zone but referenced from the Audio zone's Title Group Domain—is quite restricted compared to VTs on a DVD-Video disc. Parental control and seamless branching are not supported. Also, VTS PGCs in DVD-Audio do not include pre-commands or post-commands, and the use of dummy PGCs is not supported. Thus VTs on DVD-Audio discs are less capable of complex interactivity than those on DVD-Video discs.

Two additional notes about video on DVD-Audio discs:

- Because the contents of the VMG menu—including any motion video therein—are never seen by a DVD-Audio player, the AMG menu and VMG menu might each have a completely different look and feel, including the use of different motion video (if any).
- It is possible to create a DVD-Audio disc that plays the same content on either a DVD-Video player or a video-capable DVD-Audio player (Audio-with-Video or Universal). AOTT Audio Titles are not required. In this case, all content (other than menus) would be AVTT/AOTTs and AVTTs with PCM audio. This content would be accessed through AMGM on a video-capable DVD-Audio player and VMGM on a DVD-Video player. To ensure that this material would also play on an Audio-only player, the titles would have to be AVTT/AOTTs.

DVD-Audio Title Example

To see how the audio and video options for DVD-Audio work in practice, consider how a typical CD release with a dozen songs might be organized if it were released as a DVD-Audio disc (see Figure 4 on page 216). Group 1 on the DVD might be the same material in the same sequence as the CD release, but in multi-channel surround. Each song could be accompanied by a sequential slideshow allowing the listener to navigate to different index points in the song by clicking on lyrics. Group 2 could be a dance version of the album, with ballads removed and longer club mixes substituted for several of the other songs. Group 3 might be a collection of three or four music videos produced for the album, each enabled for playback on audio-only players. And Group 4 might be a mini-documentary including interviews with the band (with playback on audio-only players disabled.)




Figure 4: *DVD-Audio title with four groups*

Taking advantage of the disc's DVD-Others zone could open up even more creative, promotional, and commercial opportunities when the title is placed in a DVD-ROM drive. The disc might contain a custom browser programmed to take the user to the artist's section of the label's Web site, where the user might hear song samples from the band's other albums, order albums and other merchandise online, check a tour schedule and order tickets, or send fan e-mail to the band. Combining the flexibility of the DVD-Audio format with the underlying capabilities of DVD-ROM, the possibilities for DVD-Audio discs become nearly unlimited.