Intel® High Definition Audio Specification

Document Change Notification

Date: December 26, 2007

Change Identification: **DCN No: HDA-034-A2**
Document Revision: Intel® High Definition Audio 1.0

This document discloses changes to the Intel® High Definition Audio Specification and all information contained herein is provided under the terms of the "AZALIA SPECIFICATION DEVELOPMENT AGREEMENT" also known as Intel® High Definition Audio Specification Developer Agreement, and all the terms of such agreement, including the confidentiality provisions, shall apply to this disclosure.

Please note that DCN HDA-034-A2 replaces HDA-034-A. Minor modifications were made between HDA-034-A and HDA-034-A2. Intel recommends vendors to refer to HDA-034-A2 DCN. HDA-034-A will no longer be available for access via the Intel website.

Title: HDMI Content Protection and Multi-Channel Support

Brief description of the functional changes proposed:

Provides definition for a standardized method for discovery of HDMI support, and expands on providing metadata discoverability and video audio sync. Intent is to add support for HDMI as standalone capability and provide architectural definitions for proper HDCP and Multi-channel support on DRM solutions. This DCN lists changes that must be made to the HD Audio subsystem in order to support HDMI.

The Intel® High Definition Audio Specification defines a complete codec architecture that is fully discoverable and configurable so as to allow a software driver to control all typical operations of any codec. While this architectural objective is immediately intended for audio codecs, it is intended that such a standard software driver model not be precluded for modems and other codec types, such as those that support HDMI.

This DCN provides a set of additions and changes to the Intel® High Definition Audio specification that defines how a codec that supports a connection to a HDMI sink should be constructed. A reader must refer to the specification for the general requirements for an HD Audio codec.

Current Definition:

The specification only defines an HDMI function group and an HDMI value for the default pin configurations location field.

Proposed Definition:

Add capabilities to support HDMI meta-data, content protection, and multi-channel playback on the HD Audio interface.

New sections in High Definition Audio Specification v1.0:

- 7.3.3.14.1 Intrinsic Unsolicited Responses
- 7.3.3.14.2 Non-Intrinsic Unsolicited Response - Content Protection
- 7.3.3.34 HDMI EDID Like Data (ELD) Data
- 7.3.3.34.1 EDID-Like Data (ELD)
- 7.3.3.35 Converter Channel Count
- 7.3.3.36 HDMI Data Island Packet – Size info (HDMI DIP – Size)
- 7.3.3.37 HDMI Data Island Packet – Index (HDMI DIP – Index)
- 7.3.3.38 HDMI Data Island Packet – Data (HDMI DIP – Data)
- 7.3.3.39 HDMI Data Island Packet – Transmit-Control (HDMI DIP – XmitCtrl)
- 7.3.3.40 Content Protection Control (CP_CONTROL)
- 7.3.3.41 Converter Channel to HDMI Slot Mapping

3.7 Codec Verb and Response Structures

The codec verb structure is entirely opaque to the controller and link, and all fields, including the address, are only interpreted by the codec.

The controller generated (outbound) Verb format is shown in Figure 5.

Bits 31 : 28	27 : 20	19:0
Codec Address	Verb	Payload

Figure 5. Verb Format

Solicited Responses from codecs are returned by the codec in response to a command Verb. All 32 bits of the Solicited Responses are opaque to the controller and link.

The Solicited Response format is shown in Figure 6.

31:0
Response

Figure 6. Solicited Response Format

Unsolicited responses are sent by the codec independently of any software request. The 6-bit "Tag" field is opaque to the controller and used by software to distinguish what codec subunit generated the Unsolicited Response. The 5-bit "Sub Tag" field is also opaque to the controller and used by software to distinguish what widget subunit generated the Unsolicited Response (i.e. presence detect, content protection, etc.). The 20 bits of vendor specific contents can be used to provide extra contextual information to software regarding the event that generated the Unsolicited Response.

Bits 31:26	Bits 25:21	Bits 20:0
Tag	Sub Tag	Vendor Specific Contents

Figure 7. Unsolicited Response Format

7.3.3.13 Pin Widget Control

Pin Widget Control controls several aspects of the Pin Widget.

Command Options:

Table 81. Enable VRef

	Verb ID	Payload (8 Bits)	Response (32 Bits)
Get	F07h	0	Bits 31:8 are 0 Bits 7:0 are PinCntl
Set	707h	Bits 7:0 are PinCntl	0

PinCntl format:

7	6	5	4:3	2:0
H-Phn Enable	Out Enable	In Enable	Rsvd	VRefEn

Figure 63. PinCntl Format

H-Phn Enable disables/enables a low impedance amplifier associated with the output. The value 1 enables the amp. Enabling a non-existent amp is ignored. For digital pin widgets, including HDMI, this control has no function.

Out Enable allows the output path of the Pin Widget to be shut off. The value 1 enables the path. Enabling a non-existent amp is ignored. For a digital HDMI pin widget, disabling the output will cause the samples to no longer be sent to the HDMI Sink device. Other signaling may continue, such as clock recovery and other Info Frame packets.

In Enable allows the input path of the Pin Widget to be shut off. The value 1 enables the path.

VRefEn: Voltage Reference Enable controls the VRef signal(s) associated with the Pin Widget. If more than one of the bits in the VRef[7:0] field of the Pin Capabilities parameter (Section TO DO) are non-zero, then this control allows the signal level to be selected. For digital pin widgets, including HDMI, this control has no function.

The VRefEn field encoding selects one of the possible states for the VRef signal(s). If the value written to this control does not correspond to a supported value as defined in the Pin Capabilities parameter, the control must either retain the previous value or take the value of 000, which will put the control in a Hi-Z state and prevent damage to any attached components.

Table 82 enumerates the possible values for VRefEn which correlate to the values identified in the Pin Capabilities parameter.

Table 82. VRefEn Values

VRefEn Encoding	VREF Signal Level
000b	Hi-Z
001b	50%
010b	Ground (0 V)
011b	*Reserved*
100b	80%
101b	100%
110b-111b	*Reserved*

Applies to:

- Pin Complex, both Digital (HDMI) and analog

7.3.3.14 Unsolicited Responses

The **Unsolicited Response** control determines whether the node is enabled to send an unsolicited response.

This control is only available for nodes which support Unsolicited Responses, as declared in the Function Group Type parameter and the Audio Widget Capabilities parameter. The node should be queried to determine if it supports unsolicited responses before getting or setting this control.

Command Options:

Table 83. Unsolicited Responses Control

	Verb ID	Payload (8 bits)	Response (32 bits)
Get	F08h	0	Bits 31:8 are 0 EnableUnsol is bits 7:0
Set	708h	EnableUnsol is bits 7:0	0

7	6	5:0
Enable unsolicited response	Reserved	Tag

Figure 64. EnableUnSol Format

Enable Unsolicited Response controls the actual generation of Unsolicited Responses. If Enable is a 1, Unsolicited Responses may be generated. This is the only control for intrinsic unsolicited responses, such as Presence Detect. For unsolicited responses making use of sub tags, the enabling/disabling is controlled by other verbs where sub tag is provided, e.g. content protection control verb.

Note that this controls the global generation of unsolicited response for this node and if it is set to disable then unsolicited responses for any intrinsic or any sub tag type Unsolicited Responses are disabled.

Tag is a 6-bit value which is opaque to the codec and is used by software to determine what codec node generated the unsolicited response. The value programmed into the Tag field is returned in the top 6 bits

(31:26) of every Unsolicited Response generated by this node. Software can differentiate between the various unsolicited responses it receives from a given node based on the sub tag. Sub tag is provided as part of the unsolicited response body in bits 25:21 (See section 3.7).

Applies to:

- All Nodes capable of generating Unsolicited Responses.

7.3.3.14.1 Intrinsic Unsolicited Responses

Intrinsic unsolicited responses are generated as a result of asynchronous H/W events such as presence detect and hot plug (represented by ELD valid) events. These responses have a predefined sub tag of 0 and are enabled by the Unsolicited Responses Control verb.

The Unsolicited response for intrinsic events is defined as:

Bits 31:26	Bits 25:21	Bits 20:2	1	Bit 0
Tag	Sub tag	Reserved	ELDV	PD

Figure 65. Intrinsic Unsolicited Response Format

Data structure:

Table 84. Intrinsic Unsolicited Response Fields

Sub tag	This field has a predefined value of 0 for intrinsic messages
PD	Presence Detect: When this bit is set, sense measurement has changed on the pin widget and software can use the pin sense control verb to determine the current pin sense data state For analog pin widgets, this UR means that Presence Detect or Impedance has changed on the pin widget. For digital pin widgets, including HDMI pin widgets, this means that presence detect (and optionally ELD valid bit) has changed
ELDV	ELD Data valid: When this bit is set, new ELD data has been written. NOTE: Generation of this UR is gated by the Presence Detect bit, meaning that this UR can only be generated if the Presence Detect(PD) bit is already SET

The codec adheres to the following rules for generating unsolicited responses triggered by intrinsic events:

- If a second UR is generated while the previous UR is still waiting to be sent then only the new UR is sent.
- In the case of a monitor plug-in event, the ELD content is populated before the presence detect bit has been set. Setting of the ELD valid bit in this case will not generate an (extra) UR because it is gated by PD. In such cases when the PD bit is also set, an Unsolicited Response is generated that contains both PD and ELDV bits.

7.3.3.14.2 Non-Intrinsic Unsolicited Response - Content Protection

Non-Intrinsic unsolicited responses are generated as a result of synchronous events such as a Content Protection state change. The sub tag for such an unsolicited response is defined as part of the

corresponding enable command, such as Content Protection control verb and, in general, a non-zero sub tag enables the generation of unsolicited responses.

The Unsolicited response for non-intrinsic events is defined as:

Bits 31:26	Bits 25:21	Bits 20:0
Tag	Sub Tag	Command Specific

Figure 66. Generic Non-Intrinsic Unsolicited Response Format

The Unsolicited Response for content protection events is described as:

Bits 31:26	Bits 25:21	Bits 20:2	1	Bit 0
Tag	Sub Tag	Reserved	CP_State	CP_Ready

Figure 67. Content Protection Non-Intrinsic Unsolicited Response Format

The codec is enabled to generate content protection unsolicited responses only when:

a. Unsolicited responses are enabled via “Unsolicited Response Control Verb”

b. Set CP command has been sent with a non-zero sub tag

If the above two conditions are met then CP unsolicited responses are enabled and the codec *may* generate an unsolicited response *upon* either completion of CP state change request *or* upon CP state change on the link

At the hardware level, following rules govern the generation of unsolicited response for CP non-intrinsic events (assuming that codec is allowed to generate UR by meeting above two conditions):

1. UR for state change can only be generated when READY = 1 (This is consistent with the definition of READY bit, viz. CP status returned in GET CP Status verb is valid only is READY = 1. if an UR is generated while READY =0, then the status read in this case is not applicable and UR doesn’t count)

2. An UR is sent when READY bit transitions from 0 → 1 irrespective of the current CP state on the link

Above rules guarantee that only one UR is generated upon any of the above two conditions. Consider following cases:

1. **When Audio CP state change request is completed with success:** When Audio requests a CP state change, READY is cleared. Later CP state is changed at the link. Note that this doesn’t generate any UR because READY = 0 at this time. After setting CP state appropriately READY bit is set again (READY 0 → 1) and this causes generation of an UR that contains the current “changed” CP state.
2. **When Audio CP state change request is completed without success:** When Audio requests a CP state change, READY is set to 0 (READY 1 → 0). This request is ignored (by GFX driver), i.e. link CP state doesn’t change. Gfx subsystem sets the READY bit to 1 (READY 0 → 1). This causes an UR that contains the current “unchanged” CP state.
3. **CP state is changed asynchronously:** Depending on the h/w implementation, If READY stays set during the transition then change of CP state causes an UR. If the h/w clears the READY bit during the CP state change processing, then UR is generated when READY is set to 1 again.

7.3.3.15 Pin Sense

The **Pin Sense** control returns the Presence Detect status, EDID-Like Data (ELD) Valid and the impedance measurement of the device attached to the pin.

Some codecs may require that the impedance measurement be triggered by software; in that case, sending the Execute command will cause the impedance measurement to begin. The “Presence Detect” bit will always be accurate if that functionality is supported by the widget.

Note that the Pin Complex Widget may support the generation of an Unsolicited Response to indicate that the Sense Measurement (either the Presence Detect or the Impedance) value has changed, the generation of which implies that the measurement is complete.

Command Options:

Table 85. Pin Sense

	Verb ID	Payload (8 Bits)	Response (32 Bits)
Get	F09h	0	**For digital pin widget:** Bit 31: Presence Detect Bit 30: ELD Valid Bit 29:0: Reserved
			For analog pin widget: Bit 31: Presence Detect Bit 30:0: Impedance Value
Execute	709h	**For Analog pin widget:** Right Chnl: bit 0 Rsvd: bits 1:7	0
		For Digital pin widget: Not Applicable	

Presence Detect is a bit indicating the state of the Presence Detect capability. A 1 indicates that there is “something” plugged into the jack associated with the Pin Complex. This bit will only be valid if the widget has Presence Detect capability as indicated by the “Presence Detect Capable” bit of the Pin Capabilities parameter.

EDID-Like Data (ELD) Valid is a bit indicating the state of the ELD memory. When the contents are valid ELD is set to 1 and cleared to zero when not valid.

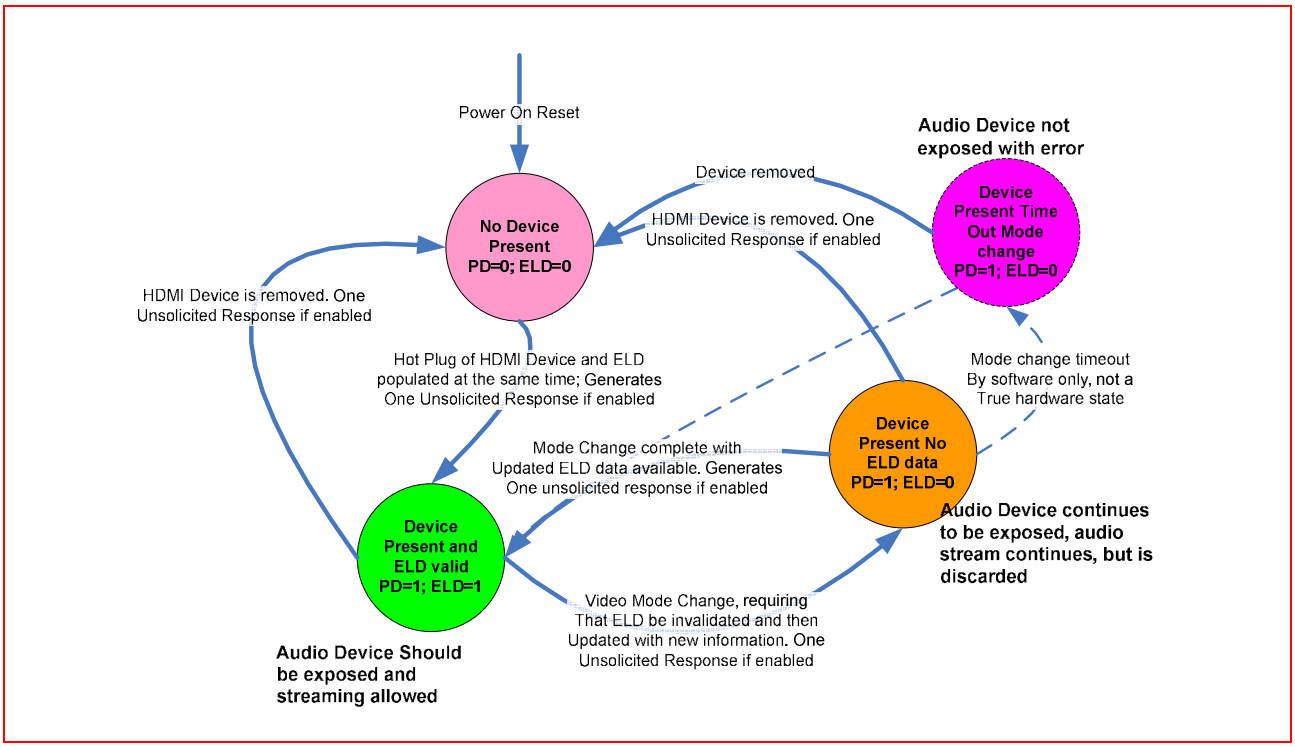

Figure 68: Presence Detect and ELD valid unsolicited responses flow for HDMI codecs

Impedance: returns the measured impedance of the widget. A returned value of 0x7FFF,FFFF (all 1's) indicates that a valid sense reading is not available, or the sense measurement is busy if it has been recently triggered. This field is only valid if the widget has Sense capability as indicated by the "Impedance Sense Capable" bit of the Pin Capabilities parameter.

Note that impedance is not valid for Digital pin widgets.

Right Chnl: Normally impedance sensing is done on the left channel or "tip" of the connector. However, Pin Widgets may optionally support sensing on the right channel or "ring" of the connector. When this bit is 1, the impedance value is taken on the right channel if the Pin Widget supports this; if not supported, this bit is ignored. When this bit is 0, the left channel is sensed.

Note that Right Chnl is not valid for Digital pin widgets.

Applies to:

- Pin Complex including Digital, HDMI and Analog

7.3.3.34 HDMI EDID-Like Data (ELD) Data

The audio software for the HDMI codec will need information about the audio capabilities of an attached HDMI sink device. This information is stored in the HDMI sink device's EDID. Typically, the EDID flows through a graphics adapter to graphics software, so the graphics adapter HW will not have knowledge of the EDID contents.

To that end, a new mechanism is defined here for passing the HDMI sink device's audio EDID information from the graphics software to the audio software. The data payload containing the audio information will be known as EDID-Like Data or ELD and will contain a subset of the HDMI sink device's EDID information. The size and contents of the ELD buffer will be determined by the HDMI Audio codec manufacturer.

The ELD information will be valid if the HDMI sink is attached and powered on and the ELD Valid bit is set. The Pin Widget that is associated with this HDMI widget will report if the device is attached and that the ELD memory is populated and valid by reporting Presence Detect of 1 and ELD Valid of 1 to a Pin Sense control command. As with the Presence Detect bit, the changes to the ELD Valid bit can also result in the generation of unsolicited responses.

Command Options:

Table 110. HDMI ELD Data

	Verb ID	Payload (8 Bits)	Response (32 Bits)
Get	F2Fh	Bits 7:0 Offset into ELD memory	Bit 31 ELD Valid indication Bits 30:8 Reserved Bits 7:0 Data: ELD data byte at specified offset into the ELD memory

Response Structure:

31	30:8	7:0
ELD Valid	*Reserved*	ELD Byte from memory

Figure 72. ELD Data Response Format

ELD Valid is a bit that indicates to software that the byte being returned is not valid.

ELD Byte [7:0] is the byte of configuration data specified by offset. For a non-existent ELD location, GET returns a value of 0. Note that the byte index will auto-increment after a Get command is completed.

Applies to:

- HDMI Pin Complex

7.3.3.35 Converter Channel Count

The *Converter Channel Count* control is used by software to program the number of channels in the incoming stream that the converter must render.

Command Options:

	Verb ID	Payload (8 Bits)	Response (32 Bits)
Get	F2Dh	0	Converter Channel Count 7:0
Set	72Dh	Converter Channel Count 7:0	

Figure 73. Converter Channel Count

The Converter Channel Count control is used to specify the number of active channels in the audio stream. It is used in conjunction with the Channel value set in the Converter Stream, Channel control to specify which channels in an incoming audio stream are to be decoded by the codec. Converter Channel Count is 0-indexed as is Channel value in the Converter Stream, Channel control.

For the example below, allow S to be the Channel value in the Converter Stream, Channel control, and C to be the Converter Channel Count. Assume two channels intended for stereo playback are destined for a codec in a stream with greater than two total channels.

Assuming the first channel to be decoded is stream position 3:

S = 2

For Stereo, 2 streams are required:

C = 1

The codec would then output a stereo stream using the third and fourth channels in the incoming stream.

Applies to:

- HDMI Output Converter

7.3.3.36 HDMI Data Island Packet – Size info (HDMI DIP – Size)

The *HDMI DIP-Size* control is used to get the sizes of the various HDMI packet buffers in the HW.

Command Options:

Table 111. HDMI DIP SIZE

	Verb ID	Payload (8 Bits)			Response (32 Bits)
Get	F2Eh	Bits 7:4 Rsvd	Bit 3 ELD buffer size	Bits 2:0 (PI) Packet Index	Bits 7:0 - 0 based size of buffer implemented for HDMI Data Island Packet PI or ELD Bits 31:8 – ZERO

Data Structure:

When Bit 3 (ELD buffer size) is set to 0:

PI Value	Definition
0h	*Audio Infoframe*
1h	*GP1* – General Purpose 1
2h	*GP1* – General Purpose 2
3h	*GP2* – General Purpose 3
4h	*GP3* – General Purpose 4
5h	*GP3* – General Purpose 5
6h	*GP3* – General Purpose 6
7h	*GP3* – General Purpose 7

Figure 74. DIP Size Packet Index

When Bit 3 (ELD Buffer size) is set to 1 then PI value is don't care and ELD buffer size is returned.

When audio data is transmitted over the HDMI link, associated control information is transmitted as "Data Island Packets" over the link. Some of the packet types applicable to audio are ACP packets (for content protection), ISRC1/2 (for content info), Audio infoframe etc. For a detailed list of control packets and respective formats, please refer to the HDMI specification. Please note that Audio Sample Packets are also DIPs like the others discussed above. This document uses DIP to refer to non-audio sample packets unless otherwise stated.

The HDMI specification defines a data island packet with a header of 4 bytes (3 bytes content + 1 byte ECC) and packet body of 32 bytes (28 bytes content and 4 bytes ECC). Note that the ECC bytes are not present in the DIP content populated by software and are hardware generated.

The present definition of data island packets does not utilize all 28 bytes in the body, giving h/w designers an opportunity to optimize by implementing smaller buffers for each packet type.

The HDMI DIP-size verbs allows software to query the size of the HDMI DIP buffers supported by the underlying hardware and make an informed decision based on the compatibility of software and h/w capabilities.

Note that an HDMI transmitter must support at least 4 DIP packet buffers (one audio infoframe and three general purpose buffers)

Applies to:

- HDMI Pin Complex

7.3.3.37 HDMI Data Island Packet – Index (HDMI DIP – index)

The *HDMI DIP – index* control sets the packet buffer index and data byte index within that packet buffer.

Command Options:

Table 112. HDMI DIP Index

	Verb ID	Payload (8 Bits)		Response (32 Bits)
Get	F30h	0		Bits 7:5 – Currently set Packet Index Bits 4:0 – Current Byte Index pointer in this packet
Set	730h	Bits 7:5 (PI) Packet Index	Bits 4:0 Byte Index (0 based byte location)	0

Packet Index (PI) selects the target Data Island Packet buffer for subsequent DIP-Data and DIP-XmitCtrl control verbs

Data Structure:

PI Value	Definition
0h	*Audio Infoframe*
1h	*GP1* – General Purpose 1
2h	*GP1* – General Purpose 2
3h	*GP2* – General Purpose 3
4h	*GP3* – General Purpose 4
5h	*GP3* – General Purpose 5
6h	*GP3* – General Purpose 6
7h	*GP3* – General Purpose 7

Figure 75. DIP INDEX packet Index

Byte Index sets the target byte offset within targeted packet buffer that is accessed in DIP-Data control verb

Software **must** ensure that *packet buffer index* and *byte index* are set to correct values before attempting to access data in the buffers or change the transmit controls. Once the *packet index* and *byte index* are set, any subsequent accesses (read/write) automatically increment the byte-index by one byte after the operation. Upon reaching the maximum possible value (32 for data island packets), it wraps around to 0.

Applies to:

- HDMI Pin Complex

7.3.3.38 HDMI Data Island Packet – Data (HDMI DIP – Data)

The *HDMI DIP – data* control accesses (Read/Write) a byte in the packet buffer previously set using *DIP-Index* control.

Command Options:

Table 113. HDMI DIP Data

	Verb ID	Payload (8 Bits)	Response (32 Bits)
Get	F31h	0	Bits 7:0 – Data byte located at target byte index in the target packet buffer
Set	731h	Bits 7:0 – Data byte to be written in the target byte index in the target packet buffer	0

Index of byte accessed via this control is determined by current byte-index value for this packet buffer. This index is set to a new value upon receiving *DIP-Index* control verb and automatically incremented after a *DIP-Data* operation. Therefore first *DIP-Data* control after *DIP-Index* control verb targets the byte determined by the *byte-index* supplied as parameter in *DIP-Index* control verb and subsequent accesses (read/write) automatically target the next byte in that buffer.

Upon reaching the maximum value (32 for data island packets), the byte index automatically wraps around to 0.

It is software's responsibility to ensure that packet index and byte index are set to correct values before it starts sending *DIP-Data* control verbs for any packet buffer.

Note that the ECC bytes are not present in the DIP content populated by DIP-Data verb, these bytes are hardware generated.

Applies to:

- HDMI Pin Complex

7.3.3.39 HDMI Data Island Packet – Transmit-Control (HDMI DIP – XmitCtrl)

The *HDMI DIP – XmitCtrl* control sets the transmission controls for the currently indexed packet buffer.

It is software's responsibility to ensure that the infoframe index is set to point to appropriate buffer and the buffer content are set to correct values before it sets the transmission control. Once transmission control is set to anything but disabled, h/w transmits the contents of indexed buffer at frequency determined by transmission control.

Command Options:

Table 114. HDMI DIP XmitCtrl

	Verb ID	Payload (8 Bits)	Response (32 Bits)
Get	F32h	0	Bits 7:6 – Transmit control for currently indexed packet buffer Bits 5:0 - Reserved
Set	732h	Bits 7:6 – Transmit control for currently indexed packet buffer Bits 5:0 - Reserved	0

Data Structure:

XmitControl Value	Definition
00	Disable Transmission
01	Reserved
10	Transmit once and then disable
11	Transmit at best effort

Figure 76. DIP XmitCtrl Value

Applies to:

- HDMI Pin Complex

7.3.3.40 Content Protection Control (CP_CONTROL)

Content protection control verb is used to set the state of content protection on the audio port. This control is only valid for pin widgets that have “CP Caps” bit set in the Audio widget capabilities response.

Command Options:

Table 115. Protection Control

	Verb ID	Payload (8 Bits)	Response (32 Bits)
Get	F33h	0	Bits 31:10 – Reserved Bit 9 – CES - Current encryption state Bit 8 – READY bit Bits 7:3 – UR sub tag for CP state Bit 2 - Reserved Bits 1:0 – Current CP Request State
Set	733h	Bits 7:3 – UR sub tag for CP state Bit 2 - Reserved Bits 1:0 –Requested CP State	0

CES Value	Definition
0	*Encryption OFF*
1	*Encryption ON*

Figure 77. Current Encryption State (CES)1

[1] Note that Current Encryption State is always the true state of encryption on the hardware and is not buffered or software settable.

READY Value	Definition
0	*H/W is NOT ready or in a state to accept CP control verbs. Audio driver should not send CP control verbs in this state*
1	*H/W is in a state to accept CP control verbs to potentially change encryption state on the audio port*

Figure 78. Ready Indication

CP State	Definition
00	*DON'T care – State can be anything*
01	*Reserved*
10	*Protection OFF – Protection state must be NO-encryption to facilitate operations like recording*
11	*Protection ON – Protection state must be ON to facilitate rendering of protected content*

Figure 79. Content Protection (CP) State

Content Protection control is used by audio driver to request setting of content protection state on the audio link. The CP state can be set to ON, OFF or DON'T CARE. Audio driver sets the CP state on the link as result of a request received from elements higher in the audio stack.

In case of HDCP/HDMI, both video and audio share the same external link as well as the h/w encryption logic. To arbitrate between audio and video side requests, the video driver acts as the master. Any request to change the CP state are actually sent to the video driver which makes decision to accept or deny the requests based on various factors such as current h/w state, video stack state etc.

Typical flow to set CP state is as follows:

a. Video subsystem initializes the HDMI link, including setting up the display hot plug registers as follows:

   Set the CP_READY bit – If the unsolicited responses are enabled, this causes an UR to be generated.

b. Audio software requests setting of a CP session and the request trickles down to the audio driver.

c. Audio software uses GET_CP_CONTROL verb to check if HDMI codec HW (in this case video subsystem) is in a state to accept requests to change CP state or not. READY bit is the indicator of hardware's capability to accept CP state change commands.

d. If READY bit is SET in the response to GET_CP_Control verb, audio software sends a SET_CP_Control verb with intended CP state parameters to the audio codec.

   This verb also contains a sub tag field that would subsequently be used to identify the "notification unsolicited responses (UR)"

   Audio software may start its timeout for "H/W acting on my request" event

e. Audio hardware upon receiving this verb sends a default response (all 0's). This response identifies to audio driver that CP request was received successfully by the h/w.

Audio H/W clears the CP_READY bit – this is an indication to the audio software that h/w CP state machine has accepted the request and has started working on it.

While CP_READY = 0, the HDCP status bit is treated as invalid by audio software (as returned in the GET_CP_CONTROL verb)

While CP_READY = 0, audio software doesn't send another CP command to the audio device.

f. CP_READY bit is also used by audio driver (in polling mode) to poll for the completion of request

g. Video subsystem interacts with HDCP h/w to start the process of changing the state of encryption (if needed)

h. Once HDCP hardware is done transitioning the state (or alternatively, video subsystem has determined that it doesn't need to change the CP state), video subsystem sets the READY bit.

Note that when the CP state on the link changes, it doesn't generate any UR. The CP state change UR is gated by CP_READY bit.

i. When CP_READY bit gets set, audio h/w sends a UR, if enabled (when SET_CP_CONTROL[sub tag] != 0 and UR are enabled)

This UR is received by audio software.

j. Audio software would clear the timeout, if any set earlier

k. Audio software can now use GET_CP_STATE verb to read the status of CP on the link

l. The GET_CP_CONTROL[status] bit indicates the true "encryption" state of the HDCP hardware

The same flow can be depicted graphically as follows:

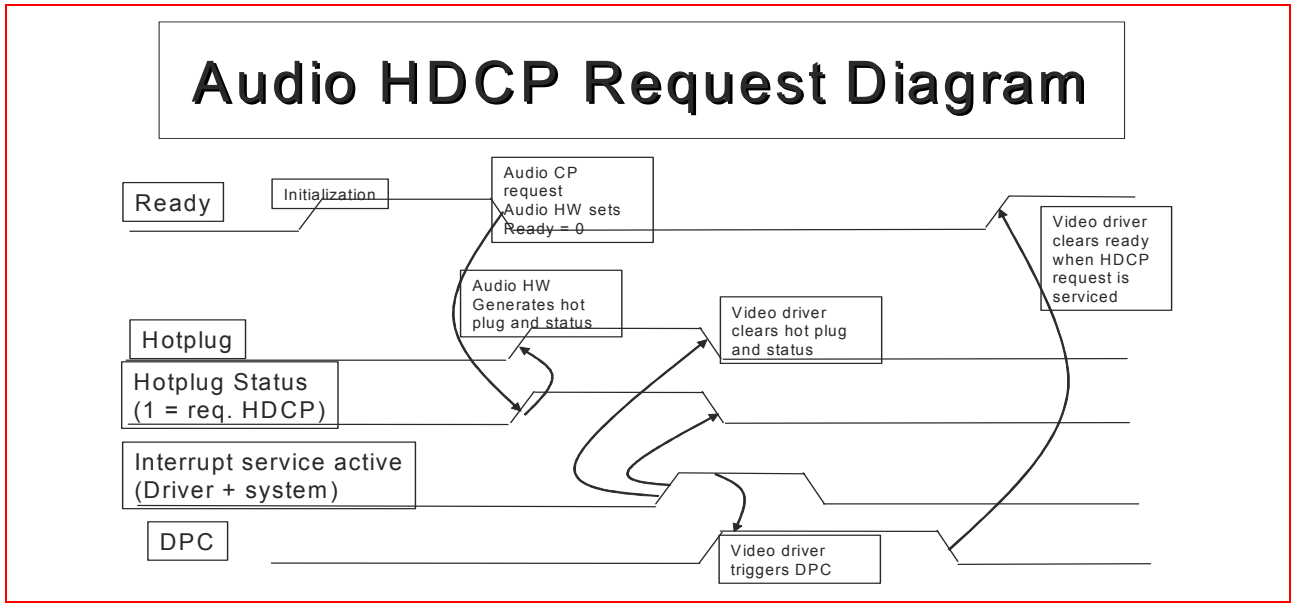

Figure 80. Audio HDCP Request Diagram

Applies to:

- HDMI Pin Complex

7.3.3.41 Converter Channel to HDMI Slot Mapping

When audio data is transmitted through the audio software stack, the stream channels are placed consecutively (tightly packed) in memory in channel order defined by the operating system, e.g. for 5.1, 16-bit content, sample for all 6 channels are placed consecutively in memory without any holes; therefore each sample consumes 12 consecutive bytes in memory. When this content is transmitted on the HD Audio link, the stream maintains its format (all 12 bytes are transmitted together in the same order as in memory).

Also, during audio content's flows through the audio software stack, speaker masks define the relative location of channels in a given stream (e.g. first 2 bytes of a 12-byte set of samples in above example are always for Front Left channel).

When audio samples are transmitted on the HDMI link, they are packetized in the form of audio sample packet (ASP). HDMI follows the definitions in CEA-861 spec for relative placement of channel data in the audio sample packet. For most cases, the mapping used in PC software is different than the mapping defined by the CEA-861 spec.

Note that in the PC environment, a common practice is to do speaker-fill – if rendered content's channel count is less than the number of speakers actually present in the system then software can map a given sample channel to multiple speakers to do "speaker fill". This allows all channels to play content. Therefore irrespective of the actual number of channels of audio content being rendered, all speakers play

some content. For example when a CD is played on a PC with 7.1 support, the audio output can appear on all 8 speakers.

In order to keep these options available for PC platforms that include HDMI, HDMI digital converters are required to support mapping of a given stream channel (converter output) to multiple HDMI channels (slots) via this verb. Note that the power on reset and function reset values for the channel-slot mapping are as described below in the notes.

"Set HDMI – Channel to slot mapping" verb is used to map an HDMI slot in audio sample packet to corresponding sample channel coming on HD Audio link into an HDMI converter.

Command Options:

Table 116. Converter Channel to HDMI Slot Mapping

	Verb ID	Payload (8 Bits)	Response (32 Bits)
Get	F34h	Bits 7:4 – Reserved Bits 3:0 – HDMI slot number	Bits 31:8 – Reserved Bits 7:4- Converter channel number Bits 3:0 – HDMI slot number
Set	734h	Bits 7:4 – Converter Channel Number Bits 3:0 – HDMI slot number	0

Note:

- The power-on reset and function-reset values for the HDMI channel to converter channel mapping are as follows:
- Converter channel 0 mapped to HDMI slot 0
- Converter channel 1 mapped to HDMI slot 1
- Converter channel 2 mapped to HDMI slot 3
- Converter channel 3 mapped to HDMI slot 2
- Converter channel 4 mapped to HDMI slot 4
- Converter channel 5 mapped to HDMI slot 5
- Converter channel 6 mapped to HDMI slot 6
- Converter channel 7 mapped to HDMI slot 7
- During setup of a stream, audio software is responsible for ensuring that all HDMI channels (slots) are mapped to corresponding converter channels (outputs).
- Specifying a mapping value of '0xF' or a value greater than Converter Channel Count means that the HDMI slot will not be driven with data (i.e. it is identified as unallocated in audio sample packet)
- The value for the HDMI slot number field in this verb is 0-based, i.e. HDMI slots are numbered from 0-7. This is different from the notation used in the CEA-861 spec where HDMI slots are

numbered from 1-8. CEA-861 HDMI slot number 1 corresponds to HDMI slot number 0 in this verb, slot number 2 in CEA-861 is noted as slot number 1 and so on.

Applies to:

- HDMI Output Converter

7.3.4.6 Audio Widget Capabilities

The Audio Capabilities control returns a set of bit fields describing the audio capabilities of the widget.

Parameter ID: 09h

Response Format:

31:24	23:20	19:16	15:13	12	11	10	9	8	7	6	5	4	3	2	1	0
Rsvd	Type	Delay	Chan count ext	CP caps	L-R Swap	Power Cntrl	Digital	Conn List	Unsol Cap-able	Proc Widget	Stripe	Format Override	Amp Param Over-ride	Out Amp Present	In Amp Present	Chan count LSB (stereo)

Figure 86. Audio Widget Capabilities Response Format

Type defines the functionality of the widget node. This is an enumerated list.

Table 118. Widget Type

Value	Type
0h	Audio Output
1h	Audio Input
2h	Audio Mixer
3h	Audio Selector
4h	Pin Complex; includes analog, digital and HDMI variants
5h	Power Widget
6h	Volume Knob Widget[2]
7h	Beep Generator Widget[3]
8h-Eh	*Reserved*
Fh	Vendor defined audio widget

[2] In the case of the Volume Knob Widget, none of the parameter bits [19:0] are used and may be omitted or set to 0. However, software assumes the capability of unsolicited responses and a connection list, as these are required by this widget type.

[3] In the case of the Beep Generator Widget, the only meaningful parameter bits are 2 ("Out Amp Present") and 3 ("Amp Param Override"). None of the other parameter bits are used and may be omitted or set to 0.

Any vendor defined widget that is enumerated hierarchically within an Audio Function Group must be identified as a vendor defined type (Fh) using this parameter.

Delay indicates the number of sample delays through the widget. This may be 0 if the delay value in the Audio Function Parameters is supplied to represent the entire path.

Chan count ext and Chan count LSB together these 4 bits specify the maximum number of channels that the widget supports. The value contained in the 4 bit field is split with the 3 most significant bits contained in bits 15:13 and the least significant bit in bit 0. These bits combined form the channel count supported minus (-) 1. So if two channels (stereo) is supported then the value would be – bits 15:13=000, bit 0=1. For 8 channels the value would be – bits 15:13=011, bit 0=1.

CP Caps indicates that the widget supports "Content Protection". No indication of the type of protection is implied by this, but does require that the Copy Protection Control verb be supported. This capability is only meaningful for pin widgets.

L-R Swap indicates the capability of swapping the left and right channels through the Audio Widget. If the Audio Widget is both input and output capable (e.g., Pin Widget), then swapping must be supported for both input and output paths. Default (0) is no swap capability.

PowerCntrl indicates that the Power State control is supported on this widget. This allows finer grained power management than just at the Function Group level for widgets which support it.

Digital indicates that a widget supports a digital stream. If the bit is a 1, it is a digital widget. For an Input or Output converter, for instance, this means the widget is translating between the High Definition Audio Link and a digital format such as S/P-DIF or I2S rather than analog data.

ConnList indicates whether a connection list is present on the widget. If the bit is a 1, the Connection List Length parameter and the Connection List Entry controls should be queried to discover the input connections. This bit must be a 1 for Input Converters, Sum Widgets, and Selector Widgets. The bit may be a 0 for Output Converters if the only connection for the widget is to the High Definition Audio Link.

If **Unsol Capable** is a 1, the audio widget supports unsolicited responses. The Unsolicited Response command can be used to configure and enable Unsolicited Response generation. When this parameter is associated with a Pin Widget, then setting this bit requires that the Pin Widget generate an unsolicited response (when enabled) whenever the "Presence Detect" bit (see Section 7.3.3.15) changes state.

If **ProcWidget** is a 1, the "Processing Controls" parameter should be queried for more information about the widget's processing controls.

The **Stripe** bit defines whether the Audio Output Converter Widget supports striping, as defined in (Section TO DO). If Stripe is a 1, the Audio Output Converter Widget must support the Stripe Control verb. Stripe is only defined for Audio Output Converter Widgets; for all other widget types, this bit must be 0.

If **Format Override** is a 1, the widget contains format information, and the "Supported Formats" and "Supported PCM Bits, Rates" should be queried for the widget's format capabilities. If this bit is a 0, then the Audio Function node must contain default amplifier parameters, and that node's format related parameters should be queried to determine the format parameters. This bit is not applicable to, and is always 0 for, Pin Complex Widgets.

If **Amp Param Override** is a 1, the widget contains its own amplifier parameters. If this bit is a 0, then the Audio Function node must contain default amplifier parameters, and they should be used to define all amplifier parameters (both input and output) in this widget.

If **(In|Out) AmpPresent** is a 1, the widget contains an input or output amplifier, as indicated. The Amp Param Override bit should be examined to determine whether the widget contains default amplifier parameters or has amplifier parameters that need to be explicitly queried. The “In Amp Present” bit is only relevant for Sum Widgets, Input Converters, and Pin Complexes. The “Out Amp Present” bit is only relevant for Selector Widgets, Sum Widgets, Output Converter Widgets, and Pin Complex Widgets.

The **Stereo** bit determines if the widget is a stereo or mono widget. If the “Stereo” bit is a 1, the widget is a stereo widget.

Applies to:

- Input Converter Widget
- Output Converter Widget
- Selector Widget
- Mixer Widget
- Pin Widget

7.3.4.7 Supported PCM Size, Rates

The Supported Rates parameter returns a bit field describing the maximum bit depth and sample rate capabilities of the widget when PCM formatted data is being rendered or captured.

Parameter ID: 0Ah

::

7.3.4.7.1 HDMI LPCM CAD

The Consolidated Audio Descriptor for LPCM content describes the current sample size/rate capability for the HDMI widget. This is potentially changed to reflect the dynamic configuration of the HDMI widget. This information will be valid if the HDMI sink is attached and powered on and the ELD Valid bit is set. Whenever the ELD bit toggles to 1, this field needs to be re-read by the software stack to figure out the updated LPCM capability of the HDMI widget.

Parameter ID: 20h

Response Format:

31:30	29:28	27:24	23:20	19:18	17:14	13:10	9:8	7:4	3:0
	Bit Rate of 192 kHz			Bit Rate of 96 kHz			Bit Rate of 48 kHz		
44.1MS/ 44.1 support	24b:20b BitRate Supported	Max Channels **without** CP ON	Max Channels with CP ON	24b:20b BitRate Supported	Max Channels **without** CP ON	Max Channels with CP ON	24b:20b BitRate Supported	Max Channels **without** CP ON	Max Channels with CP ON

Figure 74.1. Audio Widget LPCM CAD Format

44..1 support bit indicate that the HDMI sink codec supports 44.1 kHz and

44.1MS support bit indicates that the HDMI sink codec supports 44.1kHz multiples

- If this bit is 0 then 88.2 and 176.4 kHz are not supported by the sink
- *If this bit is 1* and 96kHz *fs* is supported then 88.2 kHz is also supported
- *If this bit is 1* and 192kHz *fs* is supported then 176.4 is also supported

7.3.4.9 Pin Capabilities

The Pin Capabilities parameter returns a bit field describing the capabilities of the Pin Complex Widget.

Parameter ID: 0Ch

Response Format:

31:17	16	15:8	7	6	5	4	3	2	1	0
Rsvd	EAPD Capable	VRef Control	HDMI	Balanced I/O Pins	Input Capable	Output Capable	Headphone Drive Capable	Presence Detect Capable	Trigger Req'd	Impedance Sense Capable

Figure 89. Pin Capabilities Response Format

EAPD Capable indicates the codec has an EAPD pin and that this Pin Widget provides support for controlling that pin.

VRef Control[7:0] is a bit field used to indicate what voltages may be produced on the associated VRef pin(s). If all bits in the bit field are 0, then VRef generation is not supported by the Pin Complex. Also, if the Input Capable bit is a 0, then the VRef bit field has no meaning and all bits must be 0.

If the Output Capable bit and any bits in the VRef field are set, then bit 0 (Hi-Z) must also be set to indicate that the VRef signal can be turned off to support output devices.

Figure 90 describes the VRef bit field. A 1 in any position indicates that the associated signal level is supported. All values of VRef are specified as a percentage of the analog voltage rail, AVdd.

7:6	5	4	3	2	1	0
Rsvd	100%	80%	*Rsvd*	Ground	50%	Hi-Z

Figure 90. VRef Bit Field

HDMI indicates that the Pin Complex Widget supports connection to a HDMI Sink.

Balanced I/O Pins indicates that the Pin Complex Widget has balanced pins.

Input Capable indicates whether the pin complex supports input. If Input Capable is a 1, the pin is capable of input.

Output Capable indicates whether the pin complex supports output. If Output Capable is a 1, the pin is capable of output.

Headphone Drive Capable indicates that the pin has an amplifier with sufficient current drive to drive headphones. If Output Capable is a 0, then this bit has no meaning and must be 0.

Presence Detect Capable indicates whether the pin complex can perform presence detect to determine whether there is anything plugged in. Presence detect does not indicate *what* is plugged in, only that *something* is plugged in.

Trigger Required indicates whether a trigger is required for an impedance measurement (see Section 7.3.3.15).

Impedance Sense Capable indicates whether the pin complex supports impedance sense on the attached peripheral to determine what it is. More accurate (possibly sequenced) forms of peripheral discrimination may be supported independent of this capability; however, if this bit is a 1, then the codec must support at least the basic impedance test as described in Section see Section 7.3.3.15.

Applies to:

- Pin Widget

7.3.6 Required Parameter and Control Support

Table 120 specifies which parameters are required (R) for each specification-defined node. It also indicates optional (o) parameters which are used to declare the presence of optional features in the associated node. A shaded square in the table indicates that the subject parameter is not applicable to the subject node type. The squares marked with (a) indicate an alternative; the parameter is required in either the Audio Function Group (AFG), to be used as a default, or else in all of the indicated widgets. If these parameters are present in the AFG, they are only needed in the individual widgets that have non-default capabilities. Some parameters are marked with an asterisk (*) for the "Vendor_Specific_Audio_Widget indicating they are not required by the specification since a vendor specific node may largely define its own parameters. If, however, the vendor specific node implements features that can be defined by an existing parameter, then using the standard parameter is preferable to defining a new one.

Table 120. Required Support for Parameters

Required Parameter Support	Parameter ID	Root Node	Audio Function Group	Modem Function Group	Vendor Defined Function Group	Audio Output Converter	Audio Input Converter	Pin Complex Widget	Mixer (SumAmp)	Selector (Mux)	Power Widget	Volume Knob	Beep Generator	HDMI Pin Widget	Vendor Defined Widget
Vendor ID	00	R													
Revision ID	02	R													
Subordinate Node Count	04	R	R	R	R										
Function Group Type	05		R	R	R										
Audio Function Group Capabilities	08		o												
Audio Widget Capabilities	09					R	R	R	R	R	R	R	R	R	R
Sample Size, Rate CAPs	0A		A			A	a								*
Stream Formats	0B		A			A	a								*
Pin Capabilities	0C							R						R	*
Input Amp Capabilities	0D		A				a	a	a	a					*
Output Amp Capabilities	12		A			a		a	a	a					*
Connection List Length	0E						R	R	R	R	R			R	*
Supported Power States	0F		R	R	o	o	o	o	o	o	R				*
Processing Capabilities	10					o	o	o		o					*
GPI/O Count	11		o	o	o										
Volume Knob Capabilities	13											R			
HDMI LPCM CAD	20		A			A	a							R	

Note that the Audio Function Group Capabilities parameter provides a default delay for the entire AFG to be used in lieu of adding specific delays listed for each widget in the Audio Widget Capabilities parameter. This is required if one or more widgets in the AFG opts to not report a correct delay in its Audio Widget Capabilities parameter; if all widgets do report an accurate delay number, the Audio Function Group Capabilities parameter is not required.

Table 121 specifies which verbs and controls are required (R) for each specification-defined node. It also indicates conditional (c) verbs which are required only if the respective optional capability is declared to be available. Another conditional verb (X) is required when the codec supports multiple **SDI** signals. A shaded square in the table indicates that the subject verb is not applicable to the subject node type. Some parameters are marked with an asterisk (*) for the “Vendor_Specific_Audio_Widget” indicating they are

not required by the specification since a vendor specific node may largely define its own verbs. If, however, the vendor specific node implements controls that can be accessed with an existing verb, then using the standard verb is preferable to defining a new one.

Table 121. Required Support for Verbs

Required Verb Support	Get Code	Set Code	Root Node	Audio Function Group	Modem Function Group	Group	Audio Output Converter	Audio Input Converter	Pin Complex Widget	Mixer (SumAmp)	Selector (Mux)	Power Widget	Volume Knob	Beep Generator	HDMI Pin Widget	Vendor Defined Widget
Get Parameter	F00		R	R	R	R	R	R	R	R	R	R	R	R	R	R
Connection Select	F01	701						c	c		c					*
Get Connection List Entry	F02							R	R	R	R	R	R		R	*
Processing State	F03	##					c	c	c		c					*
Coefficient Index	D- -	5- -					c	c	c		c					*
Processing Coefficient	C- -	4- -					c	c	c		c					*
Amplifier Gain/Mute	B- -	3- -					c	c	c	c	c			c		*
Stream Format	A- -	2- -					R	R								*
Digital Converter 1	F0D	70D					c	c								*
Digital Converter 2	F0D	70E					c	c								*
Power State	F05	705		R	R	c	c	c	c	c	c	R			c	c
Channel/Stream ID	F06	706					R	R								*
SDI Select	F04	704					X	X								*
Pin Widget Control	F07	707							R							*
Unsolicited Enable	F08	708					c	c	c	c	c	c	c		c	*
Pin Sense	F09	709							c							*
EAPD/BTL Enable	F0C	70C							c							*
All GPI Controls	F10 thru F1A	710 thru 71A		c	c											
Beep Generation Control	F0A	70A												R		
Volume Knob Control	F0F	70F											R			
Implementation ID	F20	720		R	R	R										
Implementation ID, Byte 1	F20	721		R	R	R										
Implementation ID, Byte 2	F20	722		R	R	R										
Implementation ID, Byte 3	F20	723		R	R	R										

Required Verb Support	Get Code	Set Code	Root Node	Audio Function Group	Modem Function Group	Group	Audio Output Converter	Audio Input Converter	Pin Complex Widget	Mixer (SumAmp)	Selector (Mux)	Power Widget	Volume Knob	Beep Generator	HDMI Pin Widget	Vendor Defined Widget
Config Default, Byte 0	F1C	71C							R						R	
Config Default, Byte 1	F1C	71D							R						R	
Config Default, Byte 2	F1C	71E							R						R	
Config Default, Byte 3	F1C	71F							R						R	
Stripe Control	F24	724					c									
Converter Channel Count	F2D	72D					c									
HDMI Data Island Pkt Size	F2E	72E													R	
HDMI ELD Data	F2F														R	
HDMI Data Island Pkt Index	F30	730													R	
HDMI Data Island Pkt Data	F31	731													R	
HDMI Data Island Pkt Cntrl	F32	732													R	
Content Protection Control	F33	733													c	
Cnv Chan – HDMI Slot Map	F34	734					c									
RESET		7FF		R	R	R										

Note that the Connection Select control is not required when the Connection List Length Register value is 1 for this node. In that case, there is no Connection Select control.

