



## PCI-SIG ENGINEERING CHANGE NOTICE

<b>TITLE:</b>	Protocol Multiplexing
<b>DATE:</b>	June 17, 2010
<b>AFFECTED DOCUMENT:</b>	PCI Express 3.0 Revision 0.9
<b>SPONSOR:</b>	AMD, HP

### **Part I**

#### **1. Summary of the Functional Changes**

This involves a minor upward compatible change in Chapter 3, Chapter 4 and a new Appendix T.

#### **2. Benefits as a Result of the Changes**

Multiple protocols can now share a PCIe Link.

#### **3. Assessment of the Impact**

No impact to systems that do not support Protocol Multiplexing. When multiplexing is enabled on a Link, PCIe traffic on that Link is unaffected. The mechanism used to identify Multiplexed Protocol Packets does not impact PCIe protocol efficiency. Protocol Multiplexing may be supported on any Port type.

#### **4. Analysis of the Hardware Implications**

Protocol Multiplexing support is optional normative. There is no impact on hardware that does not support it.


#### **5. Analysis of the Software Implications**

Feature is optional normative. There is no impact to software that does not enable the feature.



Modify Figure 3-13 in Section 3.5.2.1 as follows:

Editor: in Figure 3-13, change Symbol +0, Bits 7:4 so it has Bits 7:4 containing 0000b.



OM13787

**Figure 3-13: TLP Following Application of [TLP](#) Sequence Number and Reserved Bits**

Modify section 3.5.2.1 as follows:

### 3.5.2.1. LCRC and Sequence Number Rules (TLP Transmitter)

...


□ The following timer is used:

- **REPLAY\_TIMER** - Counts time that determines when a replay is required, according to the following rules:
  - ◆ Started at the last Symbol of any TLP transmission or retransmission, if not already running
  - ◆ For each replay, reset and restart REPLAY\_TIMER when sending the last Symbol of the first TLP to be retransmitted
  - ◆ Restarts for each Ack DLLP received while there are unacknowledged TLPs outstanding, if, and only if, the received Ack DLLP acknowledges some TLP in the retry buffer
    - Note: This ensures that REPLAY\_TIMER is reset only when forward progress is being made
  - ◆ Reset and hold until restart conditions are met for each Nak received (except during a replay) or when the REPLAY\_TIMER expires
  - ◆ Not advanced during Link retraining (holds its value when the LTSSM is in the Recovery or Configuration state). Refer to Sections 4.2.5.3 and 4.2.5.4.
  - ◆ If Protocol Multiplexing is supported, optionally not advanced during the reception of PMUX Packets (see Appendix T).
  - ◆ Resets and holds when there are no outstanding unacknowledged TLPs

...

Modify Figure 4-7 in section 4.2.1.2 as follows:

Editor: In Figure 4-7, change the text “Reserved bits and Sequence Number added by Data Link Layer” to “0000b and TLP Sequence Number added by Data Link Layer”:



OM13796

Figure 4-7: Framed TLP on a x1 Link

Modify section 4.2.2.3.1 as follows:

#### 4.2.2.3.1 Framing Tokens

...

Note that, for TLPs, the Data Link Layer prepends 4 Reserved bits (0000b) to the TLP Sequence Number field before it calculates the LCRC. These Reserved bits are not explicitly transmitted when using 128b/130b encoding, and Receivers assume that the 4 bits received are 0000b when calculating the LCRC.

Packets containing a TLP Length field that is greater than 1535 are PMUX Packets. For such packets, the actual packet length is computed differently, the TLP Sequence Number field in the STP Token contains other information, and the Link CRC is computed using different rules. See Appendix T for details.

Packets containing a TLP Length field that is between 1152 and 1535 (inclusive) are reserved for future standardization.

...

Modify section 4.2.2.3.3 as follows:

#### 4.2.2.3.3. Receiver Framing Requirements

...

- Receivers may optionally check whether the TLP Length field has a value of 0. If checked, receiving a TLP Length field of 0 is a Framing Error.
- Receivers may optionally check whether the TLP Length field has a value between 1152 and 1535 (inclusive). If checked, receiving such a TLP Length field is a Framing Error.
- Receivers on Ports that do not support Protocol Multiplexing may optionally check whether the TLP Length field has a value greater than 1535. If checked, receiving such a TLP Length field is a Framing Error.
- Receivers on Ports that support Protocol Multiplexing, shall process STP Tokens with a TLP Length field that is greater than 1535 as the start of a PMUX Packet as defined in Appendix T.

...

*Add Appendix T:*


*(Editor: Please assign the appropriate Appendix number)*

## Appendix T Protocol Multiplexing

The Protocol Multiplexing mechanism provides a standard mechanism to transport non-PCIe protocols across a PCIe Link. The mechanism supports the multiplexing of PMUX Packets and TLPs onto a single PCIe Link.

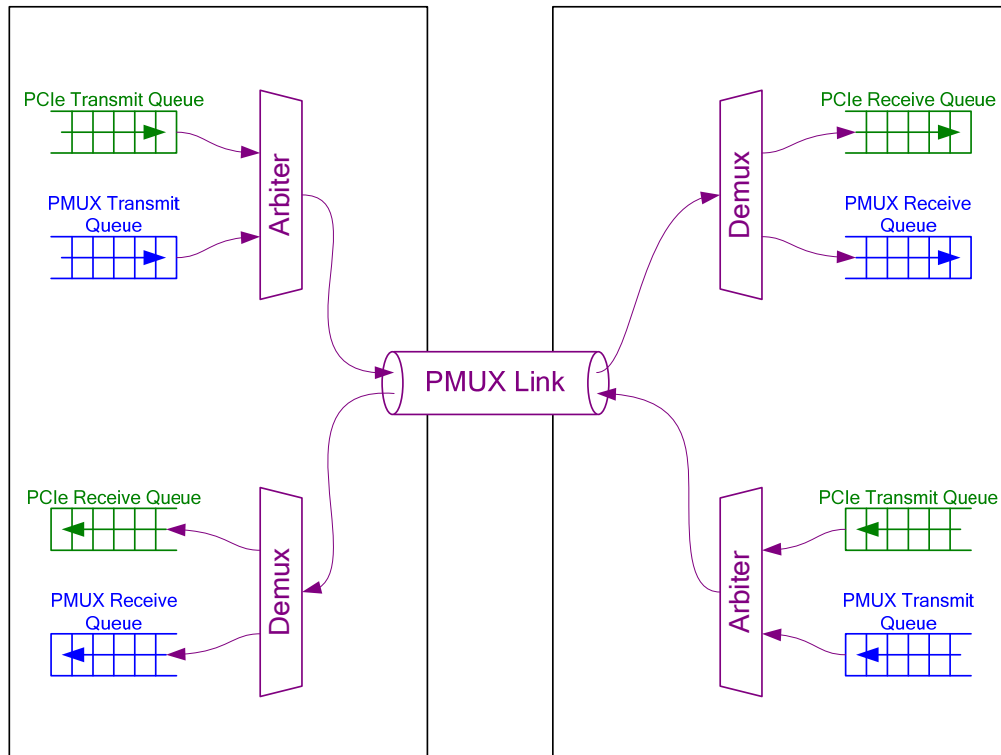
An example system topology using Protocol Multiplexing is shown in Figure T-1. In this example, the Link may operate in a two modes:

- PCIe Link. Protocol Multiplexing is disabled.
- PMUX Link. Protocol Multiplexing is enabled. Both TLPs and PMUX Packets are used in a coordinated fashion. PMUX Packets may be used to support additional protocols efficiently.



**Figure T-1: Device and Processor connected using a PMUX Link**

A PMUX Link is shown in Figure T-2. Arbitration and encapsulation occurs between the transmit queues and the Link. Demultiplexing and decapsulation occurs between the Link and the various receive queues. Packets are sent from transmit queues to the corresponding receive queues. Packets are identified as either PMUX Packets or TLPs.



**Figure T-2: PMUX Link**


Important attributes of the Protocol Multiplexing mechanism are:

- Protocol Multiplexing support is optional normative.
- Protocol Multiplexing has no impact on PCIe components that do not support it.
- Protocol Multiplexing has no impact on PCIe TLPs and DLLPs, even when it is enabled.
- A Link may be used for both TLPs and PMUX Packets at the same time.
- Protocol Multiplexing does not consume or interfere with PCIe resources (sequence numbers, credits, etc.). PMUX Packets use distinct resources associated with the specific multiplexed protocol.
- Protocol Multiplexing is disabled by default and is enabled by software. PMUX Packets must not be sent until enabled by software. PMUX Packets received at Ports that support Protocol Multiplexing are ignored until Protocol Multiplexing is enabled by software.
- Protocol Multiplexing is selectable on a per-Link basis. Protocol Multiplexing may be used on any collection of Links in a system.
- A PMUX Link may support up to 4 simultaneously active PMUX Channels. Software configures the protocol used on each PMUX Channel.
- PMUX Packets contain an LCRC. This is used to provide data resiliency in a similar fashion as PCIe TLPs.

- ❑ PMUX Packets do not use the ACK/NAK mechanism of PCIe. Multiplexed protocol specific acknowledgement mechanisms can be used to provide reliable delivery when needed.
- ❑ PMUX Packets do not contain a TLP Sequence Number. Instead, they contain a 12 bit PMUX Packet Metadata field that is available for multiplexed protocol specific use.
- ❑ PMUX Packet transmitters must contain some arbitration/QoS mechanism for scheduling sending of PMUX Packets, TLPs and DLLPs, however, the mechanism used is outside the scope of this specification.
- ❑ The Protocol Multiplexing mechanism does not define any addressing or routing mechanism for PMUX Packets.

PMUX Packets are similar to PCIe TLPs. The PMUX Packet Flow Through is shown in Figure T-3. The PCIe Packet Flow Through is shown in Figure 1-5. Changes from PCIe Packet Flow Through are:

- ❑ PMUX Packets use a protocol specific PMUX Protocol Layer instead of the PCIe Transaction Layer.
- ❑ PMUX Packets use a simplified Data Link Layer. The packet integrity portion of the Data Link Layer is mostly unchanged (LCRC computation uses a different seed value). The reliability and flow control aspects of the Data Link Layer are removed (the TLP Sequence Number field is repurposed as PMUX Packet Metadata).
- ❑ The Physical Layer is slightly modified to provide a mechanism to identify PMUX Packets.



**Figure T-3: PMUX Packet Flow Through the Layers**

Note: Figure 1-5 is reproduced here for reference and is not changed by this ECN.  
 (Editor: Please change the earlier reference to Figure 1-5 to link to the actual figure and remove this one when incorporating this ECN.)




Figure 1-5: PCIe Packet Flow Through the Layers

## T.1 Protocol Multiplexing Interactions with PCIe

Table T-1 and Table T-2 describe interactions between Protocol Multiplexing and PCIe. Table T-1 describes how a PCIe attributes affect Protocol Multiplexing. Table T-2 describes how Protocol Multiplexing features affect PCIe attributes.

Table T-1: PCIe Attribute Impact on Protocol Multiplexing

PCI Express Attribute	Impact on Protocol Multiplexing
Link Speed	<p>All PMUX Channels are disabled when the Current Link Speed corresponds to a speed that is not supported by Protocol Multiplexing (see Section T.5.4). A PMUX Channel may be disabled when the Current Link Speed corresponds to a speed that is not supported by the associated protocol.<sup>1</sup></p> <p>Link speed can change either explicitly due to a change in the Target Link Speed field or automatically due to an autonomous Link speed change (see <i>PCI Express Base Specification</i> Section 6.11).</p> <p>The PMUX Protocol Layer is permitted to influence the mechanism used by a component to determine when it requests an autonomous Link speed change. In addition, setting Hardware Autonomous Speed Disable at each end of the Link will prevent certain autonomous Link speed changes (see Section 7.8.19).</p>

<sup>1</sup> The mechanism software uses to determine what Link Speeds are supported by a protocol is outside the scope of this specification.

PCI Express Attribute	Impact on Protocol Multiplexing
	The PMUX Protocol Layer may be notified of the change.
Link Width	A PMUX Channel may be disabled when the Link Width corresponds to a width that is not supported by the associated protocol. <sup>2</sup> The PMUX Protocol Layer is permitted to influence the mechanism used by a component to determine when it requests a Link Width change. The PMUX Protocol Layer may be notified of the change.
FLR initiated	All PMUX Channels on a Link are disabled if an FLR is directed to the Upstream Port's Function 0. No PMUX Channels are affected if an FLR is directed to any other Function.
DL_Down	All PMUX Channels on a Link are disabled.
Hot Reset	All PMUX Channels on a Link are disabled.
PERST#	All PMUX Channels on a Link are disabled.
TLP Replay	No effect on Protocol Multiplexing.
DLLP Lost	No effect on Protocol Multiplexing.
TLP Prefix	No effect on Protocol Multiplexing.
Locked Transactions	No effect on Protocol Multiplexing (including no effect on any protocol specific forwarding of PMUX Packets).
AtomicOp Transactions	No effect on Protocol Multiplexing.
Multicast Transactions	No effect on Protocol Multiplexing.
Access Control Services (ACS)	No effect on Protocol Multiplexing.
Alternative Routing ID-Interpretation (ARI)	No effect on Protocol Multiplexing.
TLP Processing Hints (TPH)	No effect on Protocol Multiplexing.
Virtual Channels	No effect on Protocol Multiplexing. PCIe Links remain capable of supporting a full complement of VCs.
Internal Error	Corrected or Uncorrectable Internal Errors in the PMUX Protocol Layer may be reported as PCIe Internal Errors
L0s Link Power State	Protocol Multiplexing tracks the Link state. The PMUX Protocol Layer may request the Link transition back to L0.
L1 Link Power State	Protocol Multiplexing tracks the Link state. The PMUX Protocol Layer may request the Link transition back to L0.
Disabled LTSSM State	Disabling a Link also disables all PMUX

<sup>2</sup> The mechanism software uses to determine what Link Widths are supported by a protocol is outside the scope of this specification.

PCI Express Attribute	Impact on Protocol Multiplexing
	Channels on the Link.
Loopback LTSSM State	Entering Loopback state disables all PMUX Channels on the Link.
Recovery LTSSM State	No effect on Protocol Multiplexing. The PMUX Protocol Layer may be notified.
Receiver or Framing Error	The error is reported to the PMUX Protocol Layer to indicate that data might have been lost. This can be used to initiate protocol specific error recovery mechanisms. The Error is reported to software using PCIe Mechanisms.
Lane Reversal	No effect on Protocol Multiplexing. Support for Lane Reversal remains optional.
Polarity Inversion	No effect on Protocol Multiplexing.
Crosslink	No effect on Protocol Multiplexing. Support for Crosslink remains optional. If supported, the PMUX Protocol Layer may be notified of the outcome of the Crosslink Upstream / Downstream negotiation.
Lane assignment rules	Placement and frequency rules for STP Symbols and STP Tokens are not changed (see Section 4.2.1.2 and Section 4.2.2.3.2). These rules apply identically to PCIe TLPs and PMUX Packets.
PCI Power Management Power State	All PMUX Channels on a Link are disabled if the Upstream Port's Function 0 is sent to non-D0 state.
Dynamic Power Allocation (DPA)	No effect on Protocol Multiplexing. The PMUX Protocol Layer is notified of the change and may participate in the power reduction. PCIe power management includes any power used by the PMUX Protocol Layer.
PCI Power Management Power Consumed / Power Dissipated / Aux_Current	Power required by the PMUX Protocol Layer is included in the PCI structures.
Power Budgeting	Power required by the PMUX Protocol Layer is included in the PCIe structures.
Slot Power Limit	Slot Power Limit includes power available to PMUX Protocol Layer.
ASPM L0s Entry Condition	<p>The definition of Idle is extended to include:</p> <ul style="list-style-type: none"> <li><input type="checkbox"/> No pending PMUX Packets to transmit over the Link.</li> <li><input type="checkbox"/> For PMUX Channels that use protocol specific Flow Control, no credits are available to send PMUX Packets in that PMUX Channel.</li> </ul>
ASPM L1 Entry Condition	A Link may not enter L1 if PMUX Packets are pending or scheduled to be transmitted.
ASPM L0s/L1 Exit Conditions	A Link may be directed to exit L0s or L1 if a

<b>PCI Express Attribute</b>	<b>Impact on Protocol Multiplexing</b>
	component needs to transmit a PMUX Packet. Routing of PMUX Packets through routing elements is outside the scope of this specification; the associated L0s/L1 exit rules are also unspecified.
Bus Renumbering	No effect on Protocol Multiplexing.
Hot Plug	No direct effect on Protocol Multiplexing. Note Hot Plug events indirectly affect Data Link State which, in turn affects Protocol Multiplexing.
TLP Sequence Number	No effect on Protocol Multiplexing. PMUX Packets do not consume TLP Sequence Numbers.
PCIe Flow Control	No effect on Protocol Multiplexing. PMUX Packets do not consume PCIe Flow Control credits. Flow Control Update DLLPs must be sent as required by PCIe.
Error Reporting	No direct effect on Protocol Multiplexing. The PMUX Protocol Layer may be notified when an error is signaled or when an error message is received.
LCRC Errors in TLPs	No effect on Protocol Multiplexing.
Nullified TLPs	No effect on Protocol Multiplexing.
VC Arbitration	No effect on Protocol Multiplexing. Arbitration within PCIe is unaffected by Protocol Multiplexing.
Port Arbitration	No effect on Protocol Multiplexing. Arbitration within PCIe is unaffected by Protocol Multiplexing.
Electrical Idle Inference	PMUX Packets count as TLPs for the purpose of inferring Electrical Idle.
MR-IOV	Protocol Multiplexing may co-exist with MR-IOV. PMUX Packets are not part of any MR-IOV Virtual Hierarchy. Protocol Multiplexing is controlled using config space in the Management VH(s).

**Table T-2: PMUX Attribute Impact on PCIe**

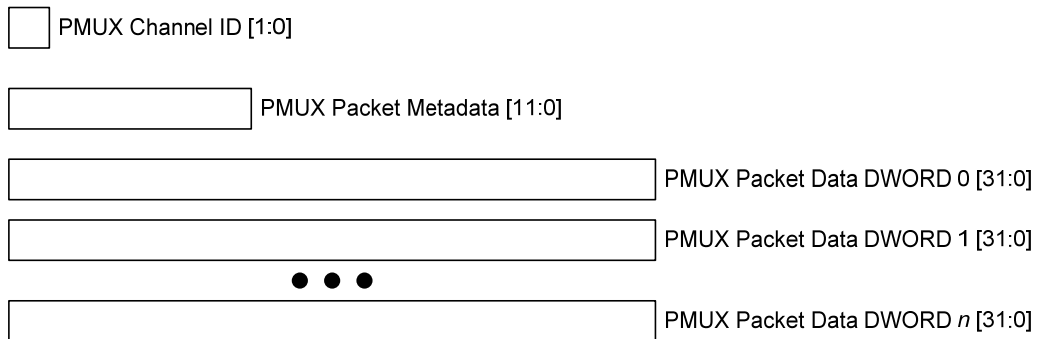
<b>PMUX Attribute</b>	<b>Impact on PCIe</b>
PMUX Protocol Error	No effect on PCIe.
LCRC Errors in PMUX Packets	No effect on PCIe. PMUX Packets with LCRC errors are discarded without triggering PCIe replay. This error is reported to the PMUX Protocol Layer and can be used to initiate protocol specific error recovery and/or error reporting mechanisms.

<b>PMUX Attribute</b>	<b>Impact on PCIe</b>
Link Unreliability	The PMUX Protocol Layer is permitted to influence the mechanism used by a component to determine if it requests an autonomous link speed change.
Nullified PMUX Packets	No effect on PCIe. It is protocol specific whether PMUX Packets within a specific PMUX Channel may be nullified. If supported, PMUX Packets are nullified in the same manner as TLPs (e.g. inverting the LCRC and signaling nullification at the Physical Layer). Receiving a nullified PMUX Packet may be reported to the PMUX Protocol Layer.
Electrical Idle Inference	PMUX Packets count as TLPs for the purpose of inferring Electrical Idle.
PMUX Protocol Layer directs LTSSM to enter Recovery	Both PCIe and the PMUX Protocol Layer are permitted to direct a transition from L0 to Recovery.
PMUX Channel Enabled / Disabled	No effect on PCIe
PMUX Packet Receiver Buffer Overflow	No effect on PCIe. This is a protocol problem within the PMUX Channel. The PMUX Transport Layer must continue to accept such packets dispose of them using protocol specific mechanisms.
Received PMUX Packet larger or smaller than supported by the associated protocol	No effect on PCIe. These are protocol problems within the PMUX Channel. The PMUX Transport Layer must accept such packets and dispose of them using protocol specific mechanisms.
Received PMUX Packet that contains more than 125 DWORDs of PMUX Packet Data	No effect on PCIe. This is an invalid PMUX Packet. The PMUX Transport Layer must accept such a packet and dispose of it. A protocol specific mechanism may be used to report the error.  Note: This situation only exists for a packet encoded using 8b10b. The TLP Length field of a packet encoded using 128b/130b cannot contain values that cause this situation.
PMUX Packet Received on disabled PMUX Channel	No effect on PCIe. No effect on any other PMUX Channel. Receivers must silently ignore such packets regardless of packet length and regardless of whether or not the packet is nullified.  PMUX Packets arriving on a disabled PMUX Channel may occur normally when software is in the process of initializing Protocol Multiplexing.
PMUX Packet Received at component that does not support Protocol Multiplexing	Software should not enable PMUX Packets unless both ends of a Link support Protocol Multiplexing.  In the 128b/130b encoding, receiving a PMUX

PMUX Attribute	Impact on PCIe
	<p>Packet by a component that does not support Protocol Multiplexing is a Framing Error (see Section 4.2.2.3.1).</p> <p>In the 8b10b encoding, the PMUX Packet LCRC is computed differently than the TLP LCRC. Receivers that do not support Protocol Multiplexing will interpret PMUX Packets as TLPs with LCRC errors and will not process them.</p>
<p>Large PMUX Packets when PCIe Max_Payload_Size is small</p>	<p>Under certain conditions, it is possible for a large PMUX Packet to trigger a premature PCIe replay. For example, this can occur when the time needed to transmit a PMUX Packet is larger than the REPLAY_TIMER (see Section 3.5.2.1).</p> <p>To avoid this issue, implementations are permitted to not advance (hold) their REPLAY_TIMER during the reception of PMUX Packets.</p> <p>Note: The PCIe REPLAY_TIMER mechanism has adequate headroom for most cases. This issue exists when (1) Max_Payload_Size is 000b, (2) PMUX Packets are larger than about 80 DWORDs, and (3) the REPLAY_TIMER is at the low end of the -0%/+100% tolerance.</p>

## T.2 PMUX Packets

A PMUX Packet contains the information shown in Figure T-4



**Figure T-4: PMUX Packet**

PMUX Channel ID is a 2 bit field that identifies which protocol is associated with a PMUX Packet. PMUX Channel ID values are between 0 and 3 (inclusive).

PMUX Packet Metadata is a 12 bit field that provides information about the PMUX Packet. Definition of this field is protocol specific and is outside the scope of this specification.

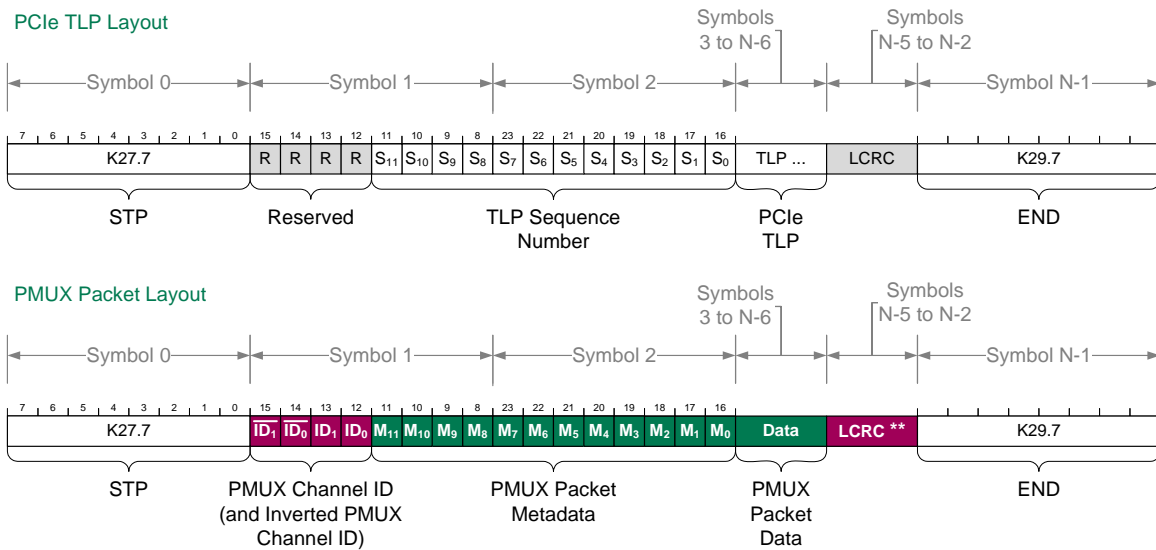
A PMUX Packet consists of between 0 and 125 DWORDs of PMUX Packet Data. Layout and usage of these DWORDs is protocol specific and is outside the scope of this specification. A PMUX Packet need not have any PMUX Packet Data and may consist only of PMUX Channel ID and PMUX Packet Metadata.

### T.3 PMUX Packet Layout

There are two layouts defined for PMUX Packets. One layout is used for 2.5 and 5.0 GT/s data rates and another layout is used for 8.0 GT/s and higher data rates. These layouts are discussed in the following sections.

#### T3.1 PMUX Packet Layout for 8b10b encoding

Figure T-5 and Table T-3 show the layout of PMUX Packets when using 8b10b encoding. For reference, the 8b10b encoding of a TLP is also shown (see Section 4.2.1.2 for the official definition). In Table T-3, items shown in *italics* are identical in PMUX Packets and TLPs.



\*\* LCRC for PMUX Packets uses a different starting seed than that used for TLPs . This avoids aliasing problems and potential errors if software misconfigures a link so that PMUX Packets are seen by existing silicon.

Figure T-5: TLP and PMUX Packet Framing (8b10b encoding)

**Table T-3: PMUX Packet Layout (8b10b encoding)**

Symbol	Field	Bit Position(s)	PMUX Packet Usage	TLP Usage
0	<i>Start TLP Indicator</i>	7:0	<i>K27.7</i>	
1	Inverted PMUX Channel ID [1:0]	7:6	Inverted (1s complement) of Symbol 1 bits 6:4.	Reserved
	PMUX Channel ID[1:0]	5:4	PMUX Channel ID.	Reserved
	PMUX Packet Metadata[11:8]	3:0	PMUX Packet Metadata[11:8]	TLP Sequence Number[11:8]
2	PMUX Packet Metadata[7:0]	7:0	PMUX Packet Metadata[7:0]	TLP Sequence Number[7:0]
3 to N-6	Packet	7:0	PMUX Packet	TLP
N-5 to N-2	LCRC	7:0	PMUX LCRC	PCIe LCRC
N-1	<i>END</i>	7:0	<i>K29.7</i>	

For PMUX Packets, symbols 1 and 2 contain PMUX Packet Metadata in the same bit positions that TLPs use for TLP Sequence Number.

The PMUX LCRC algorithm is identical to the TLP LCRC algorithm as described in Section 3.5.2 with the following modifications.

- The seed value is FB3E E248h (TLP LCRC uses FFFF FFFFh).
- The PMUX Channel ID field in Symbol 1 bits 7:4 is included in the PMUX LCRC in the same manner as the 4 reserved bits in the TLP LCRC.
- The PMUX Packet Metadata field is included in the PMUX LCRC in the same manner as the TLP Sequence Number field is included in the TLP LCRC.



## IMPLEMENTATION NOTE

### PMUX Packets at Receivers that do not support Protocol Multiplexing

The bits used for PMUX Channel ID are reserved unless Protocol Multiplexing is supported. As such, Receivers that do not support Protocol Multiplexing must ignore the PMUX Channel ID bits. If software misconfigures Protocol Multiplexing, a component that does not support Protocol Multiplexing could receive a PMUX Packet. To prevent that component from misinterpreting such a PMUX Packet as a valid TLP, the LCRC computation is changed for PMUX Packets. The result is that a valid PMUX Packet will never be misinterpreted as a valid TLP. These LCRC “errors” may trigger PCIe replay and may result in REPLAY\_NUM Rollover correctable errors being reported.



## IMPLEMENTATION NOTE

### PMUX Packet LCRC

The PMUX Channel ID field is covered by the LCRC. As such, at when using 8b10b encoding, receivers must wait until the LCRC is checked to make firm decisions based on the PMUX Channel ID value. The Inverted PMUX Channel ID can be compared against the PMUX Channel ID to make tentative decisions.

Note: The value of the LCRC associated with a given PMUX Packet is independent of the encoding used to transmit the packet.

### T.3.2 PMUX Packet Layout at 128b/130b encoding

Figure T-6 and Table T-4 show the layout of PMUX Packets when using 128b/130b encoding. For reference, the 128b/130b encoding of a TLP is also shown (see Section 4.2.2.2 for the official definition). In Table T-4, items shown in *italics* are identical in PMUX Packets and TLPs.

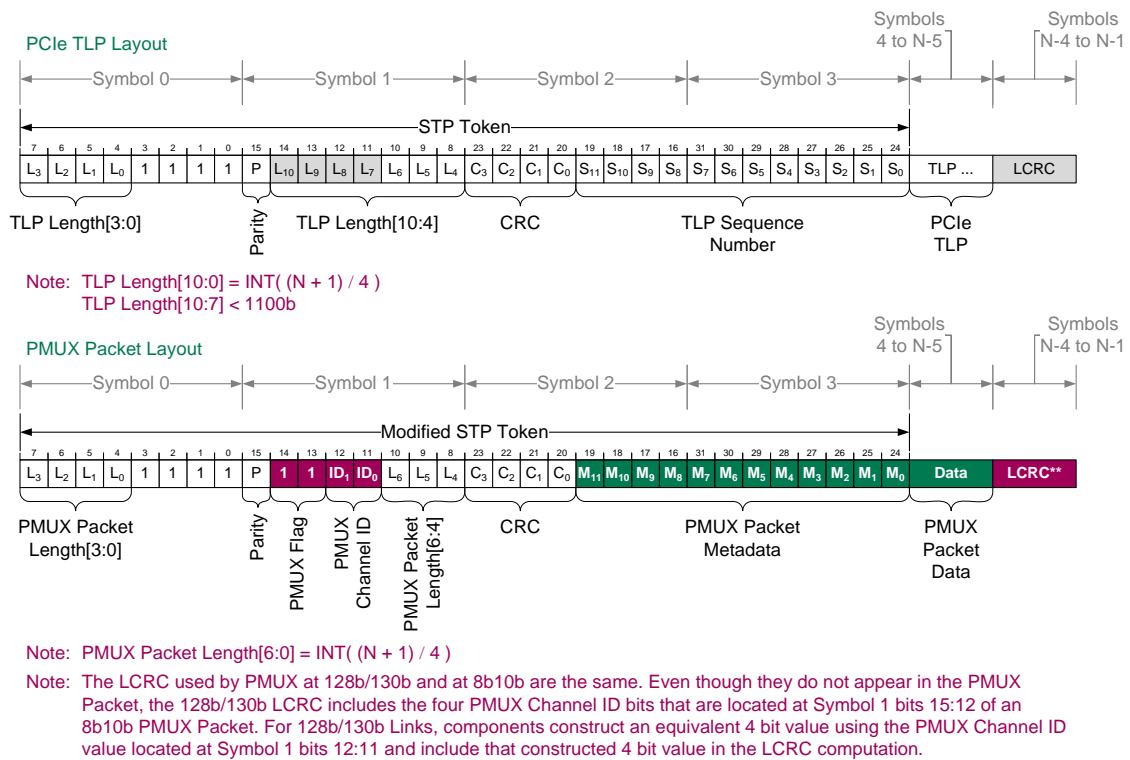


Figure T-6: TLP and PMUX Packet Framing (128b/130b encoding)

**Table T-4: PMUX Packet Layout (128b/130b encoding)**

Symbol	Field	Bit Position(s)	PMUX Packet Usage	TLP Usage
0	<i>Start TLP Indicator</i>	3:0	<i>Value of 1111b</i>	
	PMUX Packet Length[3:0]	7:4	Bits [3:0] of the PMUX Packet Length. Bit 0 is the least significant PMUX Packet Length bit.	Bits [3:0] of the TLP Length field. Bit 0 is the least significant TLP Length bit.
1	<i>Frame Parity (P)</i>	7	<i>Even parity of Symbol 0 bits [7:4], Symbol 1 bits [6:0] and Symbol 2 bits [7:4]</i>	
	PMUX Packet Indicator	6:5	Value of 11b	Bits [10:4] of the TLP Length field. Bit 10 is the most significant TLP Length bit.
	PMUX Channel ID[1:0]	4:3	PMUX Channel ID.	
	PMUX Packet Length[6:4]	2:0	Bits [6:4] of PMUX Packet Length. Bit 6 is the most significant PMUX Packet Length bit.	
2	PMUX Packet Metadata[11:8]	3:0	PMUX Packet Metadata[11:8]	TLP Sequence Number[11:8]
	<i>Frame CRC (C[3:0])</i>	7:4	<i>CRC of Symbol 0, bits [7:4] and Symbol 1 bits [6:0].</i>	
3	PMUX Packet Metadata[7:0]	7:0	PMUX Packet Metadata[7:0]	TLP Sequence Number[7:0]
4 to N-5	Packet	7:0	PMUX Packet	TLP
<i>N-4 to N-1</i>	<i>LCRC</i>	7:0	<i>LCRC</i>	

Table T-5 describes the encodings of Symbol 1 bits [6:3] in more detail. If these bits contain a value less than 1001b, the packet is a TLP and is processed as described in section 4.2.2.<sup>3</sup> If these bits contain 1001b, 1010b or 1011b, the encoding is reserved for future standardization and is processed as described in Section 4.2.2.3.3. If these bits contain a value greater than or equal to 1100b, the packet is a PMUX Packet is defined as specified in this appendix.<sup>4</sup>

<sup>3</sup> The value 1001b supports a maximum TLP Length [10:0] value of 1151 DWORDs (decimal). This will accommodate a TLP consisting of 4096 bytes of payload, 16 bytes of TLP Header, 4 bytes of TLP digest, and 480 bytes of TLP Prefix.

<sup>4</sup> The value 1100b was chosen to simplify distinguishing PMUX Packets from TLPs and from the reserved encodings.

**Table T-5: Symbol 1 bits [6:3]**

Symbol 1 bits [6:3]	Meaning
0xxx or 1000b	Packet is a TLP. Bits [6:3] are TLP Length [10:7].
1001b, 1010b, or 1011b	Encoding reserved for future standardization. Receivers detecting these encodings shall process them as described in Section 4.2.2.3.3.
1100b	Packet is a PMUX Packet. PMUX Channel ID is 0.
1101b	Packet is a PMUX Packet. PMUX Channel ID is 1.
1110b	Packet is a PMUX Packet. PMUX Channel ID is 2.
1111b	Packet is a PMUX Packet. PMUX Channel ID is 3.

For PMUX Packets, the packet length in DWORDs is contained in PMUX Packet Length [6:0]. Other than being a smaller field, PMUX Packet Length is interpreted in the same manner as TLP Length. Specifically, PMUX Packet Length also includes the framing and PMUX LCRC DWORDs (see Section 4.2.2.2).

For PMUX Packets, symbols 2 and 3 contain PMUX Packet Metadata in the same bit positions that TLPs use for TLP Sequence Number.

The PMUX LCRC algorithm is identical to the TLP LCRC algorithm as described in Section 3.5.2 with the following modifications.

- ❑ The seed value is FB3E E248h (TLP LCRC uses FFFF FFFFh).
- ❑ The PMUX Channel ID field in Symbol 1 bits 4:3 is used to compute a 4 bit value that is included in the PMUX LCRC in the same manner as the 4 reserved bits in the TLP LCRC. This 4 bit value contains the value that would be used, by the 8b10b encoding, for Symbol 1 bits 7:4. Specifically the lower 2 bits of this 4 bit value contain the PMUX Channel ID and the upper 2 bits contain the inverse (1s complement) of the PMUX Channel ID.
- ❑ The PMUX Packet Metadata field is included in the PMUX LCRC in the same manner as the TLP Sequence Number field is included in the TLP LCRC.

The Frame CRC and Frame Parity fields are computed as shown below. This is the same algorithm computed over the same bit positions as defined in Section 4.2.2.2.

$$\begin{aligned}
 C[0] &= 1b \wedge \text{PMUX\_Channel\_ID}[0] \wedge L[6] \wedge L[4] \wedge L[2] \wedge L[1] \wedge L[0] \\
 C[1] &= 1b \wedge 1b \wedge \text{PMUX\_Channel\_ID}[0] \wedge L[5] \wedge L[4] \wedge L[3] \wedge L[2] \\
 C[2] &= 1b \wedge \text{PMUX\_Channel\_ID}[1] \wedge L[6] \wedge L[4] \wedge L[3] \wedge L[2] \wedge L[1] \\
 C[3] &= \text{PMUX\_Channel\_ID}[1] \wedge \text{PMUX\_Channel\_ID}[0] \wedge L[5] \wedge L[3] \wedge L[2] \wedge L[1] \\
 &\quad \wedge L[0] \\
 P &= 1b \wedge 1b \wedge \text{PMUX\_Channel\_ID}[1] \wedge \text{PMUX\_Channel\_ID}[0] \wedge L[6] \wedge L[5] \wedge \\
 &\quad L[4] \wedge L[3] \wedge L[2] \wedge L[1] \wedge L[0] \wedge C[3] \wedge C[2] \wedge C[1] \wedge C[0]d
 \end{aligned}$$



## IMPLEMENTATION NOTE

### PMUX Channel ID and Frame CRC

When using 128b/130b encoding, the PMUX Channel ID field is covered by the Frame CRC and Frame Parity fields. As such, receivers may make decisions based on the PMUX Channel ID value as soon as the Frame CRC and Frame Parity is checked and need not wait until the PMUX LCRC is checked.

Note: The PMUX Channel ID is also covered by the LCRC. The value of the LCRC associated with a given PMUX Packet is independent of the encoding used to transmit the packet.

---

### T.4 PMUX Control

Protocol Multiplexing is disabled by default. Each PMUX Channel must be explicitly enabled by software at each end of the associated Link. Protocol Multiplexing is disabled whenever the link drops (Data Link Layer indicates DL\_Down).

A component that supports Protocol Multiplexing indicates such by the presence of the PMUX Extended Capability.

The following rules apply to components that support Protocol Multiplexing:

- PMUX Packets received in a PMUX Channel that is not enabled are silently ignored.
- PMUX Packets may not be transmitted unless the associated PMUX Channel is enabled. A PMUX Channel may also require additional, protocol specific, initialization mechanisms before PMUX Packets may be transmitted.

### T.5 PMUX Extended Capability

Figure T-7 shows the PMUX Extended Capability structure. The presence of this capability indicates that the Port supports the optional Protocol Multiplexing mechanism. This capability is optional and may be present in any Downstream Port and in Function 0 of any Upstream Port. It must not be present in non-zero Functions of Upstream Ports or in RCRBs.

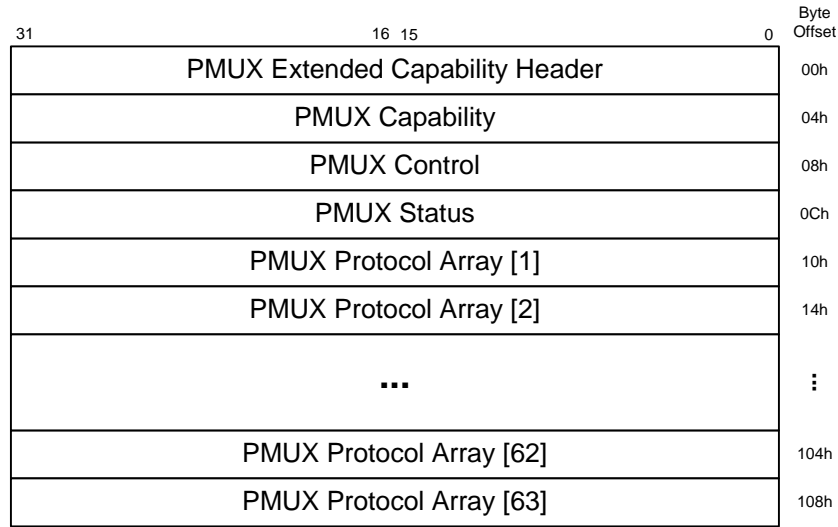
The length of the PMUX Extended Capability is determined by the PMUX Protocol Array Size field (see Section T.5.2).

This capability contains a list of the protocols supported by the Link (the PMUX Protocol Array). It also contains the mechanism software uses to enable and configure PMUX Channels. This capability must be present in both the Upstream and Downstream Ports of a Link in order for Protocol Multiplexing to be successfully enabled.

Software may enable the Upstream and Downstream Ports of a Link in either order. Software may enable multiple PMUX Channels using a single write to the PMUX Control Register.

Behavior is undefined if software enables Protocol Multiplexing in one Port and the other Port of the Link does not support Protocol Multiplexing. Behavior is also undefined if


software configures a PMUX Channel inconsistently (the same PMUX Channel in the Ports on each end of a Link configured with incompatible protocols).



**Figure T-7: PMUX Extended Capability**

T.5.1 PCI Express Extended Header (Offset 00h)

Figure T-8 details the allocation of fields in the PMUX Extended Capability header; Table T-6 provides the respective bit definitions.



**Figure T-8: PMUX Extended Capability Header**


**Table T-6: PMUX Extended Capability Header**

Bit Location	Register Description	Attributes
15:0	<b>PCI Express Extended Capability ID</b> – This field is a PCI-SIG defined ID number that indicates the nature and format of the Extended Capability. The Extended Capability ID for the PMUX Extended Capability is 001Ah.	RO
19:16	<b>Capability Version</b> – This field is a PCI-SIG defined version number that indicates the version of the Capability structure present. Must be 1h for this version of the specification.	RO

Bit Location	Register Description	Attributes
31:20	<b>Next Capability Offset</b> – This field contains the offset to the next PCI Express Capability structure or 000h if no other items exist in the linked list of capabilities. This offset is relative to the beginning of PCI compatible Configuration Space and thus must always be either 000h (for terminating the list of Capabilities) or greater than 0FFh.	RO

## T.5.2 PMUX Capability Register (Offset 04h)

Figure T-9 details the allocation of fields in the PMUX Capability register. Table T-7 provides the respective bit definitions.



**Figure T-9: PMUX Capability Register**

**Table T-7: PMUX Capability Register**

Bit Location	Register Description	Attributes
5:0	<b>PMUX Protocol Array Size</b> – Indicates the size of this Function’s PMUX Protocol Array. This field may be 0 to indicate that even though no protocols are supported, the Port will ignore all received PMUX Packets.	RO
7:6	<b>Reserved</b>	RsvdP

Bit Location	Register Description	Attributes								
15:8	<p><b>PMUX Supported Link Speeds</b> – This field indicates the Link speed(s) where Protocol Multiplexing is supported. Each bit corresponds to a Link speed. If a bit is Set, Protocol Multiplexing is supported at that Link speed. If a bit is Clear, Protocol Multiplexing is not supported at that Link speed.</p> <p>Bit definitions are:</p> <table> <tr> <td>Bit 8</td> <td>2.5 GT/s</td> </tr> <tr> <td>Bit 9</td> <td>5.0 GT/s</td> </tr> <tr> <td>Bit 10</td> <td>8.0 GT/s</td> </tr> <tr> <td>Bits 15:11</td> <td>RsvdP</td> </tr> </table> <p>At least one Link speed must be supported (i.e. the field must be non-zero). A Port may support any combination of Link speeds. For example, this field could contain the value 0000 0100b indicating that Protocol Multiplexing is only supported at 8.0 GT/s.</p> <p>This field must not indicate support for Link speeds that are not supported by the Link (see Section 7.8.18).</p> <p>Note that this field indicates the Link speeds supported by Protocol Multiplexing for the Link. The Link speeds that a particular protocol supports and the mechanism used to report that information are protocol specific.</p>	Bit 8	2.5 GT/s	Bit 9	5.0 GT/s	Bit 10	8.0 GT/s	Bits 15:11	RsvdP	RO / RsvdP
Bit 8	2.5 GT/s									
Bit 9	5.0 GT/s									
Bit 10	8.0 GT/s									
Bits 15:11	RsvdP									
31:16	<b>Reserved</b>	RsvdP								


### T.5.3 PMUX Control Register (Offset 08h)

Figure T-10 details the allocation of fields in the PMUX Control register. Table T-8 provides the respective bit definitions.

Channel  $n$  is enabled and available for use by the PMUX Protocol Layer when all of the following are true:

- The Channel  $n$  Assignment field is non-zero.
- The Channel  $n$  Assignment field is less than or equal PMUX Protocol Array Size.
- The Channel  $n$  Assignment field indicates an implemented entry in the PMUX Protocol Array (see Section T.5.5).
- All of the PMUX Channel  $n$  Disabled bits are Clear (see Section T.5.4).

Otherwise, Channel  $n$  is disabled.



**Figure T-10: PMUX Control Register**

**Table T-8: PMUX Control Register**

<b>Bit Location</b>	<b>Register Description</b>	<b>Attributes</b>
5:0	<p><b>PMUX Channel 0 Assignment</b> – This field indicates the protocol assigned to PMUX Channel 0. If the field is 0h, no protocol is assigned. If the field is non-zero, it is the index in the PMUX Protocol Array of the protocol assigned to PMUX Channel 0.</p> <p>If PMUX Protocol Array Size is less than 63 (see Section T.5.2), unused upper bits of this field may be hardwired to 0b. If PMUX Protocol Array Size is 0, this entire field may be hardwired to 0.</p> <p>This field defaults to 0h.</p>	RW
7:6	<b>Reserved</b>	RsvdP
13:8	<p><b>PMUX Channel 1 Assignment</b> – This field indicates the protocol assigned to PMUX Channel 1. If the field is 0h, no protocol is assigned. If the field is non-zero, it is the index in the PMUX Protocol Array of the protocol assigned to PMUX Channel 1.</p> <p>If PMUX Protocol Array Size is less than 63 (see Section T.5.2), unused upper bits of this field may be hardwired to 0b. If PMUX Protocol Array Size is 0, this entire field may be hardwired to 0.</p> <p>This field defaults to 0h.</p>	RW
15:14	<b>Reserved</b>	RsvdP
21:16	<p><b>PMUX Channel 2 Assignment</b> – This field indicates the protocol assigned to PMUX Channel 2. If the field is 0h, no protocol is assigned. If the field is non-zero, it is the index in the PMUX Protocol Array of the protocol assigned to PMUX Channel 2.</p> <p>If PMUX Protocol Array Size is less than 63 (see Section T.5.2), unused upper bits of this field may be hardwired to 0b. If PMUX Protocol Array Size is 0, this entire field may be hardwired to 0.</p> <p>This field defaults to 0h.</p>	RW
23:22	<b>Reserved</b>	RsvdP
29:24	<p><b>PMUX Channel 3 Assignment</b> – This field indicates the protocol assigned to PMUX Channel 3. If the field is 0h, no protocol is assigned. If the field is non-zero, it is the index in the PMUX Protocol Array of the protocol assigned to PMUX Channel 3.</p> <p>If PMUX Protocol Array Size is less than 63 (see Section T.5.2), unused upper bits of this field may be hardwired to 0b. If PMUX Protocol Array Size is 0, this entire field may be hardwired to 0.</p> <p>This field defaults to 0h.</p>	RW
31:30	<b>Reserved</b>	RsvdP


## T.5.4 PMUX Status Register (Offset 0Ch)

Figure T-11 details the allocation of fields in the PMUX Status register. Table T-9 provides the respective bit definitions.

Each channel has a set of Disabled bits. When Channel  $n$  Assignment field is non-zero, the Channel  $n$  Disabled bits reflect the error status of the channel. The following Disabled bits are defined:

- PMUX Channel  $n$  Disabled: Link Speed
- PMUX Channel  $n$  Disabled: Link Width
- PMUX Channel  $n$  Disabled: Protocol Specific

When there are multiple reasons for disabling a channel, an implementation may choose which reason(s) to report. For example, if a protocol needs bandwidth equivalent to  $x1\ 8.0\ \text{GT/s}$ , when there is inadequate bandwidth (e.g. the Link is operating at  $x1\ 5.0\ \text{GT/s}$ ,  $x1\ 2.5\ \text{GT/s}$ , or  $x2\ 2.5\ \text{GT/s}$ ) it could disable the PMUX Channel by indicating any or all of Disabled: Link Width, Disabled: Link Speed or Disabled: Protocol Specific.



**Figure T-11: PMUX Status Register**

**Table T-9: PMUX Status Register**

Bit Location	Register Description	Attributes
0	<b>PMUX Channel 0 Disabled: Link Speed</b> – If Set, Channel 0 is disabled because the Current Link Speed (Section 7.8.8) is not supported by Protocol Multiplexing or by the protocol assigned to Channel 0. This bit is 0 when no protocol is assigned to Channel 0 (i.e. Channel 0 Control field is 0h).	RO
1	<b>PMUX Channel 0 Disabled: Link Width</b> – If Set, Channel 0 is disabled because the current Link Width is not supported by the protocol assigned to Channel 0. This bit is 0 when no protocol is assigned to Channel 0 (i.e. Channel 0 Assignment field is 0h).	RO


Bit Location	Register Description	Attributes
2	<b>PMUX Channel 0 Disabled: Protocol Specific</b> – If Set, Channel 0 is disabled for protocol specific reasons. This bit is 0 when no protocol is assigned to Channel 0 (i.e. Channel 0 Assignment field is 0h).	RO
7:3	<b>Reserved</b>	RsvdZ
8	<b>PMUX Channel 1 Disabled: Link Speed</b> – If Set, Channel 1 is disabled because the Current Link Speed (Section 7.8.8) is not supported by Protocol Multiplexing or by the protocol assigned to Channel 1. This bit is 0 when no protocol is assigned to Channel 1 (i.e. Channel 1 Assignment field is 0h).	RO
9	<b>PMUX Channel 1 Disabled: Link Width</b> – If Set, Channel 1 is disabled because the current Link Width is not supported by the protocol assigned to Channel 1. This bit is 0 when no protocol is assigned to Channel 1 (i.e. Channel 1 Assignment field is 0h).	RO
10	<b>PMUX Channel 1 Disabled: Protocol Specific</b> – If Set, Channel 1 is disabled for protocol specific reasons. This bit is 0 when no protocol is assigned to Channel 1 (i.e. Channel 1 Assignment field is 0h).	RO
15:11	<b>Reserved</b>	RsvdZ
16	<b>PMUX Channel 2 Disabled: Link Speed</b> – If Set, Channel 2 is disabled because the Current Link Speed (Section 7.8.8) is not supported by Protocol Multiplexing or by the assigned protocol. This bit is 0 when no protocol is assigned to Channel 2 (i.e. Channel 2 Assignment field is 0h).	RO
17	<b>PMUX Channel 2 Disabled: Link Width</b> – If Set, Channel 2 is disabled because the current Link Width is not supported by the assigned protocol. This bit is 0 when no protocol is assigned to Channel 2 (i.e. Channel 2 Assignment field is 0h).	RO
18	<b>PMUX Channel 2 Disabled: Protocol Specific</b> – If Set, Channel 2 is disabled for protocol specific reasons. This bit is 0 when no protocol is assigned to Channel 2 (i.e. Channel 2 Assignment field is 0h).	RO
23:19	<b>Reserved</b>	RsvdZ
24	<b>PMUX Channel 3 Disabled: Link Speed</b> – If Set, Channel 3 is disabled because the Current Link Speed (Section 7.8.8) is not supported by Protocol Multiplexing or by the assigned protocol. This bit is 0 when no protocol is assigned to Channel 3 (i.e. Channel 3 Assignment field is 0h).	RO

Bit Location	Register Description	Attributes
25	<b>PMUX Channel 3 Disabled: Link Width</b> – If Set, Channel 3 is disabled because the current Link Width is not supported by the assigned protocol. This bit is 0 when no protocol is assigned to Channel 3 (i.e. Channel 3 Assignment field is 0h).	RO
26	<b>PMUX Channel 3 Disabled: Protocol Specific</b> – If Set, Channel 3 is disabled for protocol specific reasons. This bit is 0 when no protocol is assigned to Channel 3 (i.e. Channel 3 Assignment field is 0h).	RO
32:27	<b>Reserved</b>	RsvdZ

### T.5.5 PMUX Protocol Array (Offsets 10h through 48h)

The PMUX Protocol Array consists of up to 63 entries. The size of the PMUX Protocol Array is indicated by the PMUX Protocol Array Size field (see Section T.5.2).

Figure T-12 details the allocation of fields in each PMUX Protocol Array entry. Table T-10 provides the respective bit definitions.



**Figure T-12: PMUX Protocol Array Entry**

**Table T-10: PMUX Protocol Array Entry**

Bit Location	Register Description	Attributes
15:0	<b>Protocol ID</b> – In conjunction with Authority ID designates a specific protocol and the mechanism by which that protocol is mapped onto Protocol Multiplexing.	RO
31:16	<b>Authority ID</b> – Designates the authority controlling the values used in the Protocol ID field. The Authority ID field contains a Vendor ID as assigned by the PCI-SIG.	RO

The value 0000 0000h indicates an unimplemented PMUX Protocol Array entry. The PMUX Protocol Array is indexed starting at 1.

PMUX Channel *n* is enabled and configured to support the protocol associated with PMUX Protocol Array entry at index *m* when the PMUX Channel *n* Assignment field contains the value *m* (see Section T.5.3).

Entries in the PMUX Protocol Array with the Authority ID value 1 (0001h) represent protocols that are defined by the PCI-SIG.

Duplicate Entries in the PMUX Protocol Array may be used to represent multiple instances of a particular protocol. This permits software control of the mapping between PMUX Channel ID and a specific instance of a protocol.

---



## IMPLEMENTATION NOTE

### Multiple Protocol Instances

A Link may have a single PMUX Protocol assigned to multiple PMUX Channels. Each PMUX Channel is assigned to a different instance of the protocol. Each instance of a protocol corresponds to an entry in the PMUX Protocol Array.

Consider a Port that supports 2 instances of protocol X. Two entries in the PMUX Protocol Array would indicate protocol X (indexes A and B for example). To assign instance A to PMUX Channel 0 and instance B to PMUX Channel 2, place the value A in the PMUX Channel 0 Assignment field and the value B in the PMUX Channel 2 Assignment field.

---