



PCI-SIG ENGINEERING CHANGE NOTICE

TITLE:	Precision Time Measurement (PTM), Revision 1.0a
DATE:	Updated: 31 March 2013 (errata fix)
AFFECTED DOCUMENT:	PCI Express Base Specification v3.0, PCI Code and ID Assignment Specification v1.3
SPONSOR:	Intel, HP

Part I

1. Summary of the Functional Changes

Defines an optional-normative Precision Time Measurement (PTM) capability. To accomplish this, Precision Time Measurement defines a new protocol of timing measurement/synchronization messages and a new capability structure.

2. Benefits as a Result of the Changes

PTM provides several advantages including but not limited to:

Shared Time Domain Across Multiple Components: PTM enables the coordination of timing information across multiple devices with independent local timebases.

New Uses: A number of time sensitive media and server applications that are not practical when software alone is responsible for coordinating time within the platform become simple to implement using PTM.

3. Assessment of the Impact

There is no impact on devices that do not implement PTM.

4. Analysis of the Hardware Implications

PTM is an Optional Normative capability. Endpoints, Switches and Root Complexes may optionally implement PTM. The capability for precisely maintaining a local time clock may be required for some applications not currently served by PCIe hardware mechanisms.

5. Analysis of the Software Implications

System software must determine which components support PTM and appropriately enable components.

With PTM, supporting components determine the relationship between their local time and a PTM Master Time. Use of this relationship by software is implementation specific.

6. Analysis of the C&I Test Implications

No identified impact to existing C&I tests.

Part II

Detailed Description of the change

Add the following line to the Terms and Acronyms section:

Terms and Acronyms

Precision Time Measurement (PTM)

Optional capability for communicating precise timing information between components.

Modify section 2.2.8 as follows:

This document defines the following groups of Messages:

INTx Interrupt Signaling

...

OBBF Messages

PTM Messages

...

All Message Requests use the Msg Type field encoding, except for the Vendor_Defined Messages, which can use either Msg or MsgD, and the Set_Slot_Power_Limit and PTM ResponseD Messages which use MsgD.

Insert section 2.2.8.x as follows:

2.2.8.x. Precision Time Measurement (PTM) Messages

Table 2-x defines the Precision Time Measurement (PTM) Messages.

The PTM Request and PTM Response Messages must use a TLP Type of Msg, and must not include a data payload. The Length field is reserved.

- Figure 2-xi illustrates the format of the PTM Request and Response Messages.

The PTM ResponseD Message must use a TLP Type of MsgD, and must include a 64 bit PTM Master Time field in bytes 8 through 15 of the TLP header and a 1 DW data payload containing the 32 bit Propagation Delay field.

- Figure 2-xii illustrates the format of the PTM ResponseD Message.

- Refer to section 6.x.2.2 for details regarding how to populate the PTM ResponseD Message.

The Requester ID must be set to the Transmitting Port's ID.

A PTM dialog is defined as a matched pair of messages consisting of a PTM Request and the corresponding PTM Response or PTM ResponseD message.

- The PTM Messages must use the default Traffic Class designator (TC0). Receivers implementing PTM must check for violations of this rule. If a receiver determines that a TLP violates this rule, it must handle the TLP as a Malformed TLP.
 - This is a reported error associated with the Receiving Port (see Section 6.2).

Table 2-x: Precision Time Measurement Messages

<u>Name</u>	<u>TLP Type</u>	<u>Code[7:0] (b)</u>	<u>Routing r[2:0] (b)</u>	<u>Support</u>				<u>Description/Comments</u>
				<u>R</u> <u>C</u>	<u>E</u> <u>P</u>	<u>S</u> <u>W</u>	<u>B</u> <u>R</u>	
<u>PTM Request</u>	<u>Msg</u>	<u>0101 0010</u>	<u>100</u>	<u>r</u>	<u>t</u>	<u>tr</u>		<u>Initiates PTM dialog</u>
<u>PTM Response</u>	<u>Msg</u>	<u>0101 0011</u>	<u>100</u>	<u>t</u>	<u>r</u>	<u>tr</u>		<u>Completes current PTM dialog – does not carry timing information</u>
<u>PTM ResponseD</u>	<u>MsgD</u>	<u>0101 0011</u>	<u>100</u>	<u>t</u>	<u>r</u>	<u>tr</u>		<u>Completes current PTM dialog - carries timing information</u>

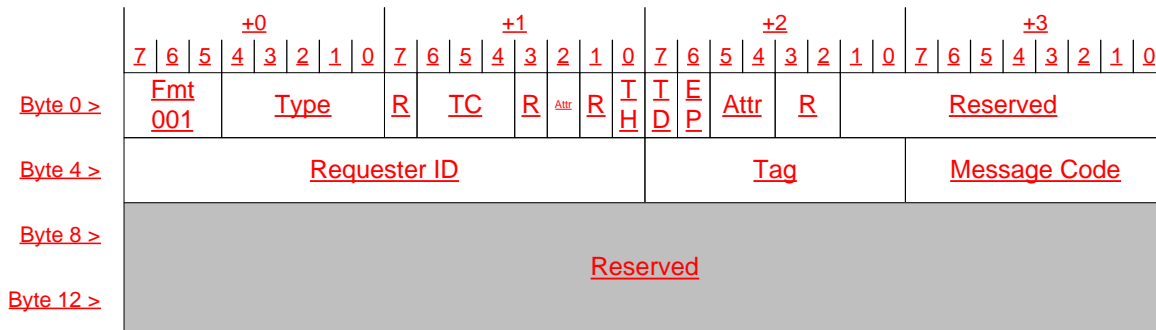


Figure 2-xi: PTM Request/Response Message

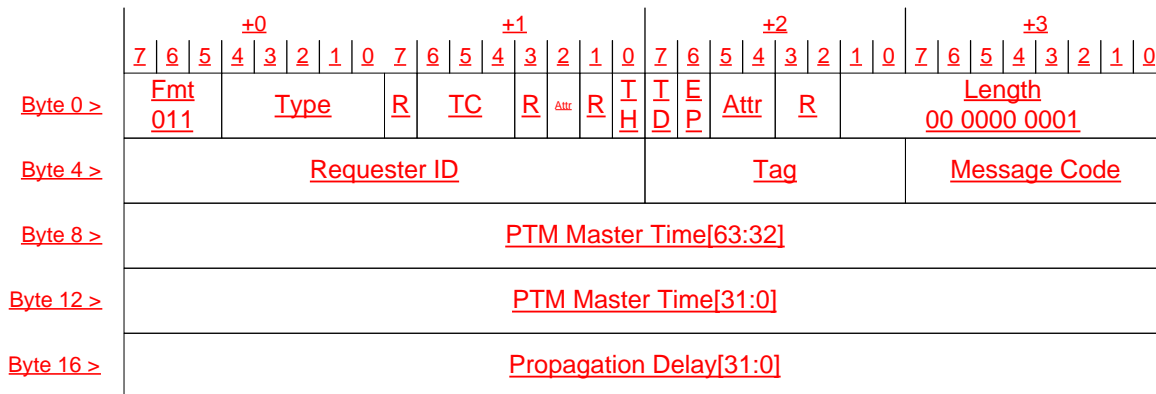


Figure 2-xii: PTM ResponseD Message (4 DW header and 1 DW payload)

Insert section 6.x as follows:

6.x Precision Time Measurement (PTM) Mechanism

6.x.1 Introduction

Precision Time Measurement (PTM) enables precise coordination of events across multiple components with independent local time clocks. Ordinarily, such precise coordination would be difficult given that individual time clocks have differing notions of the value and rate of change of time. To work around this limitation, PTM enables components to calculate the relationship between their local times and a shared PTM Master Time: an independent time domain associated with a PTM Root.

PTM defines the following:

- PTM Requester – A Function capable of using PTM as a consumer associated with an Endpoint or an Upstream Port.
- PTM Responder – A Function capable of using PTM to supply PTM Master Time associated with a Port or an RCRB.
- Time Source – A local clock associated with a PTM Responder.

□ PTM Root – The source of PTM Master Time for a PTM Hierarchy. A PTM Root must also be a Time Source and is typically also a PTM Responder.

Each PTM Root supplies PTM Master Time for a PTM Hierarchy: a set of PTM Requesters associated with a single PTM Root.

Figure 6-a. illustrates some example system topologies using PTM. These are only illustrative examples, and are not intended to imply any limits or requirements.

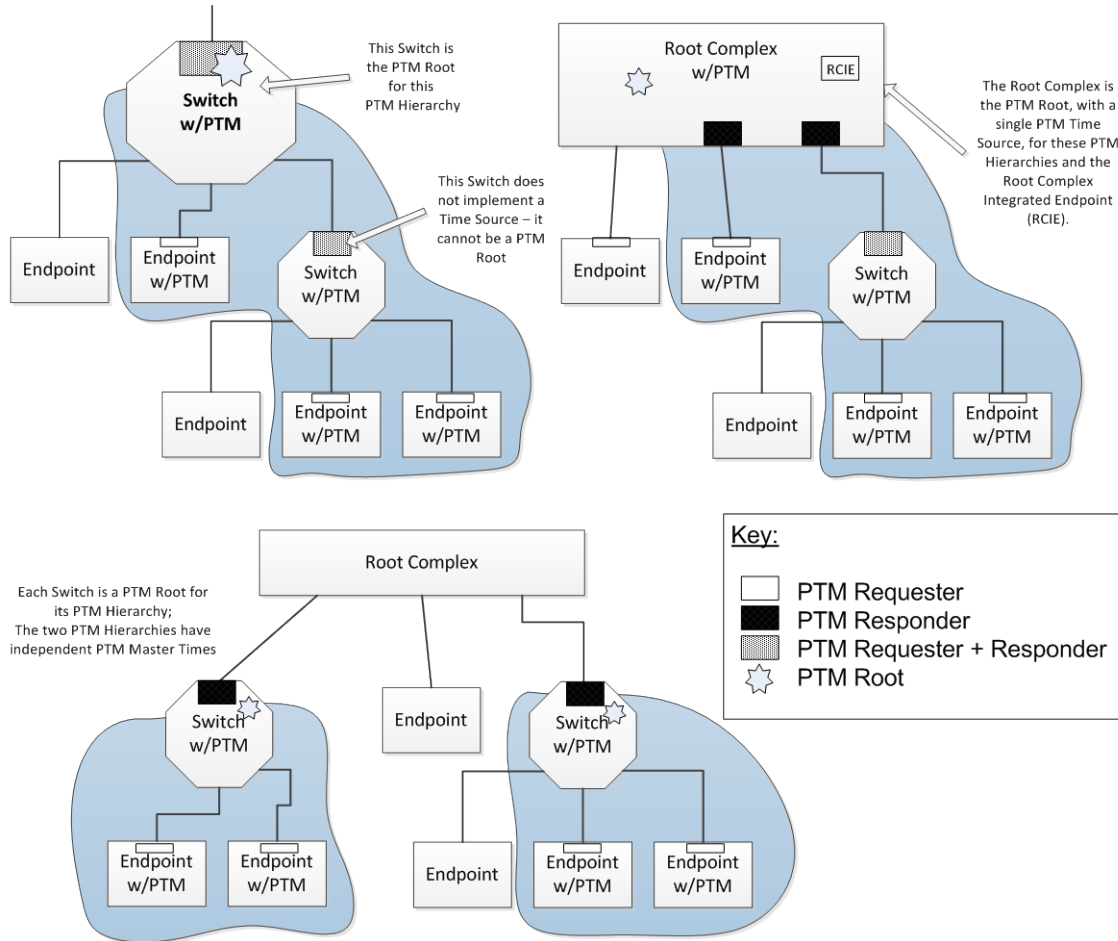


Figure 6-a: Example System Topologies using PTM

6.x.2 PTM Link Protocol

When using PTM between two components on a Link, the Upstream Port, which acts on behalf of the PTM Requester, sends PTM Requests to the Downstream Port on the same Link, which acts on behalf of the PTM Responder.

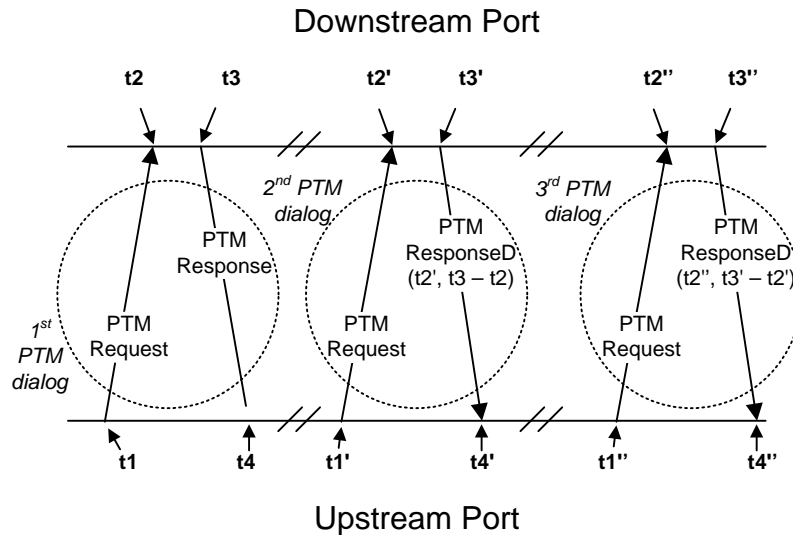


Figure 6-b: Precision Time Measurement Link Protocol

Figure 6-b illustrates the PTM link protocol. The points t1, t2, t3, and t4 in the above diagram represent timestamps captured locally by each Port as they transmit and receive PTM Messages. The component associated with each Port stores these timestamps from the 1st PTM dialog in internal registers for use in the 2nd PTM dialog, and so on for subsequent PTM dialogs.

The Upstream Port, on behalf of the PTM Requester, initiates the PTM dialog by transmitting a PTM Request message.

The Downstream Port, on behalf of the PTM Responder, has knowledge of or access (directly or indirectly) to the PTM Master Time.

During each dialog, the Downstream Port populates the PTM ResponseD message based on timestamps stored during previous PTM dialogs, as defined in Section 6.x.2.2.

Once each component has historical timestamps from the preceding dialog, the component associated with the Upstream Port can combine its timestamps with those passed in the PTM ResponseD message to calculate the PTM Master Time using the following formula:

$$\text{PTM Master Time at } t1' = t2' - \frac{((t4 - t1) - (t3 - t2))}{2}$$

The values t1, t2, t3, t4, and t2' indicate the timestamps captured during the PTM dialog as illustrated in Figure 6-b.

PTM capable components would typically record the results of these timestamp calculations, and may make them available to software via implementation specific means.

Herein, this document refers to this resultant timing information as the component's "PTM context".

For a Switch implementing PTM, the time synchronization mechanism(s) within the Switch itself are implementation-specific.



IMPLEMENTATION NOTE

PTM Theory and Operation

The timestamps captured during the PTM dialogs enable the calculation of the timing relationship between the PTM Requester and PTM Responder. The value $(t3-t2)$ measures the time consumed by the PTM Responder for a given PTM dialog. The time $(t4-t1)$ is the time from request to response. Therefore $((t4 - t1) - (t3 - t2))$ effectively gives the round trip message transit time between the two components, and that quantity divided by 2 approximates the Link delay - the time difference between $t1$ and $t2$. It is assumed that the Link transit times from PTM Requester to PTM Responder and back again are symmetric, which is typically a good assumption (see also the Implementation Note on PTM Timestamp Capture Mechanisms).

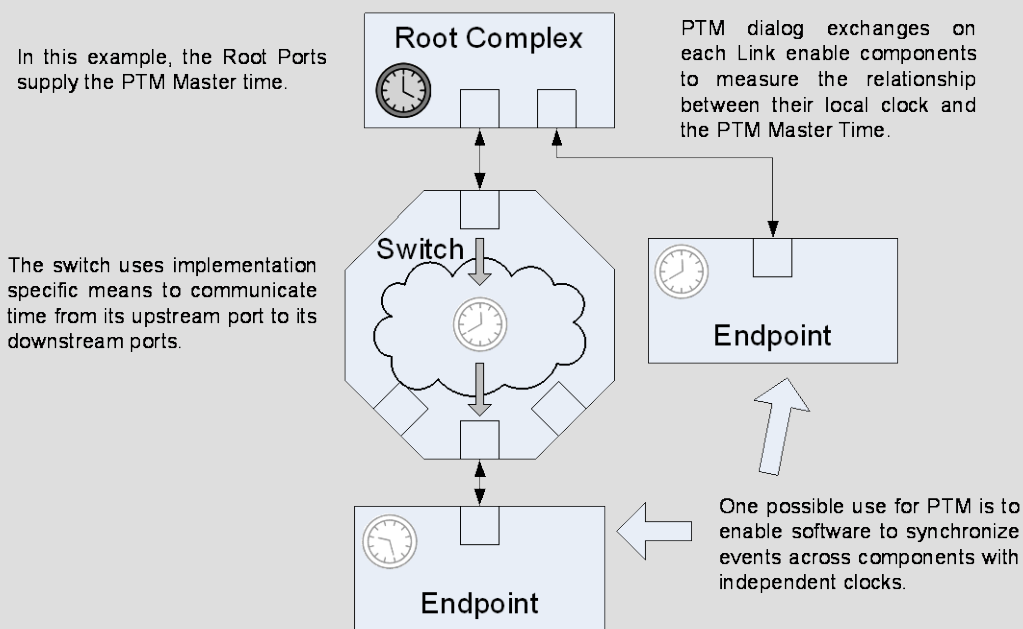


Figure 6-c: Precision Time Measurement Example

Figure 6-c illustrates a simple device hierarchy employing PTM. Each Upstream Port initiates PTM dialogs to establish the relationship between its local time and the PTM Master Time provided by the Root Port.

In this example, the Switch initiates PTM dialogs on its Upstream Port to obtain the PTM Master Time for use in fulfilling PTM Request Message received at its Downstream Ports. This Switch employs implementation specific means to communicate the PTM Master Time from its Upstream Port to its Downstream Ports.

PTM capable components can make their PTM context available for inspection by software, enabling software to translate timing information between local times and PTM Master Time. In turn, this capability enables software to coordinate events across multiple components with very fine precision.

Similarly, it is strongly recommended that platforms implementing PTM also make the PTM Master Time available to software.

6.x.2 Configuration and Operational Requirements

Software must not have the PTM Enable bit Set in the PTM Control register on a Function associated with an Upstream Port unless the associated Downstream Port on the Link already has the PTM Enable bit Set in its associated PTM Control register.

PTM support by a Function is indicated by the presence of a PTM Extended Capability structure. It is not required that all Endpoints in a hierarchy support PTM, and it is not required for software to enable PTM in all Endpoints that do support it.

If a PTM Message is received by a Port that does not support PTM, or by a Downstream Port when the PTM Enable bit is clear, the Message must be treated as an Unsupported Request. This is a reported error associated with the Receiving Port (see Section 6.2). A Properly formed PTM Response received by an Upstream Ports that support PTM, but for which the PTM Enable bit is clear, must be silently discarded.

As observed through PTM, the PTM Master Time must satisfy the following behavioral requirements:

- Time values must be monotonic, and strictly increasing.
- The perceived granularity must be no greater than the value reported in the Local Clock Granularity field of the PTM Capability register.
- The perceived time must start no later than when the PTM Root processes its first PTM Request Message.

Referring to Figure 6-b, the following rules define timestamp capture:

- A PTM Requester must update its stored t1 timestamp when transmitting a PTM Request Message, even if that transmission is a replay.
- A PTM Responder must update its stored t2 timestamp when receiving a PTM Request Message, even if received TLP is a duplicate.
- A PTM Responder must update its stored t3 timestamp when transmitting a PTM Response or ResponseD Message, even if that transmission is a replay.
- A PTM Requester must update its stored t4 timestamp when receiving a PTM Response Message, even if received TLP is a duplicate.
- Timestamps must be ~~measured~~ based on the STP Symbol or Token that frames the TLP, as if ~~observed from~~ observing the first bit of that Symbol or Token at the Port's pins.
 - Typically this will require an implementation specific adjustment to compensate for the inability to directly measure the time at the actual pins, as the time will commonly be measured at some internal point in the Rx or Tx path. The accuracy and consistency of this measurement are not bounded by this specification, but it is strongly recommended that the highest practical level of accuracy and consistency be achieved.

6.x.2.1. PTM Requester Role

- Support for the PTM Requester role is indicated by setting the PTM Requester Capable bit in the PTM Capability register.

- ❑ PTM Requesters are permitted to request PTM Master Time only when PTM is enabled. The mechanism for directing a PTM Requester to issue such a request is implementation specific.
 - Upstream Ports obtain PTM Master Time via PTM dialogs as described in section 2.2.8.x.
 - The mechanism by which Root Complex Integrated Endpoints request PTM Master Time is implementation specific.
- ❑ Once having issued a PTM Request Message, the Upstream Port must not issue another PTM Request Message prior to the receipt of a PTM Response Message, PTM ResponseD Message, Reset, or the passage of 100µs without a corresponding PTM Message from the Downstream Port.
- ❑ Upon receiving a PTM Response, the Upstream Port must wait at least 1µs before issuing another PTM Request Message.
- ❑ For multi-function devices (MFDs) containing multiple PTM Requesters, the Upstream Port associated with that MFD must issue a single PTM dialog during each PTM context refresh. PTM Requesters within the MFD maintain their individual PTM contexts using this one, Device-wide PTM dialog. The mechanism for refreshing multiple PTM contexts from one PTM dialog is implementation specific.
- ❑ It is strongly recommended that an Upstream Port invalidate its internal PTM context when the relationship between PTM Master Time and the Upstream Port's local time changes, as determined by implementation specific criteria. For example, this may occur as a result of a transition to a non-D0 state or due to accumulated PPM drift. These events are grouped under the label "Local Time Invalidation Event" in Figure 6-d.

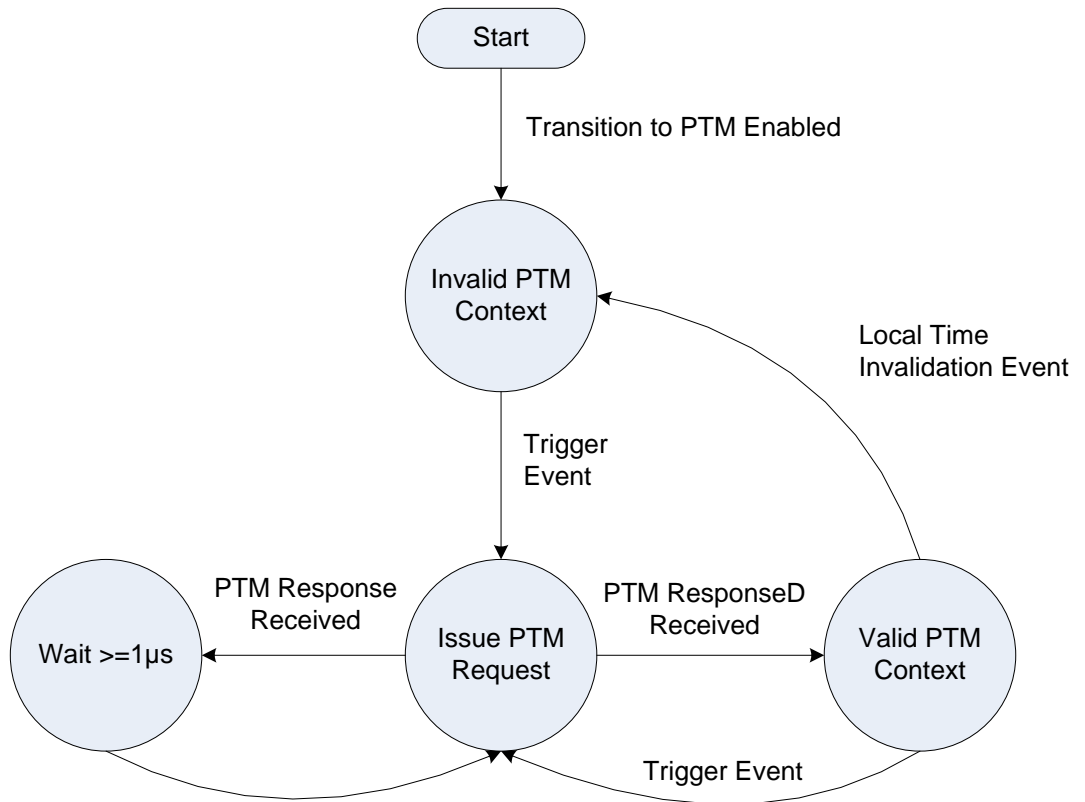


Figure 6-d: PTM Requester Operation

6.x.2.2. PTM Responder Role

- ❑ Support for the PTM Responder role is indicated by setting the PTM Responder Capable bit in the PTM Capability register.
- ❑ Switches and Root Complexes are permitted to implement the PTM Responder Role.
 - A PTM capable Switch, when enabled for PTM by setting the PTM Enable bit in the PTM Control register associated with the Switch Upstream Port, must respond to all PTM Request Messages received at any of its Downstream Ports.
 - The mechanism by which Root Complexes communicate PTM Master Time to Root Complex Integrated Endpoints is implementation specific.
- ❑ PTM Responders must populate PTM ResponseD Messages as follows (refer to Figure 6-b and the accompanying implementation note):
 - The PTM Master Time field is a 64-bit value containing the value of PTM Master Time at the receipt of the PTM Request Message for the current PTM Dialog. In Figure 6-b, for the 2nd PTM dialog, this is the PTM Master Time at time t2'.
 - The Propagation Delay field is a 32-bit value containing the interval between the receipt of the PTM Request Message and the transmission of the PTM Response Message for the previous PTM dialog. In Figure 6-b, for the 2nd PTM dialog, this is the time interval between t2 and t3 captured during the 1st PTM dialog.
 - The unit of measurement for both fields is one ns.
- ❑ Switch Downstream Ports and Root Ports acting as PTM Responders must respond to each PTM Request Message received at their Downstream Ports with either PTM Response or PTM ResponseD according to the following rules:
 - A PTM Responder must not send a PTM Response or PTM ResponseD Message without first receiving a PTM Request Message.
 - Upon receipt of a PTM Request Message, a PTM Responder must attempt to issue a PTM Response or PTM ResponseD Message within 10 μ s.
 - A PTM Responder must issue PTM Response when the Downstream Port does not have historical timestamps (t3 – t2) with which to fulfill a PTM Request Message.
 - A PTM Responder must issue PTM ResponseD when it has stored copies of the values required to populate the PTM ResponseD Message: historical timestamps (t3 –t2) and the PTM Master Time at the receipt of the most recent PTM Request Message (time t2').
 - A PTM Responder is permitted to issue PTM Response when it has stored copies of the historical timestamps (t3 – t2) but must request the PTM Master Time from elsewhere. In this case, it is permitted to issue PTM Response messages in response to PTM Request Messages while it retrieves the PTM Master Time if that retrieval is expected to take more than 10 μ s.

- The perceived granularity of the historical timestamps and PTM Master Time values transmitted by a PTM Responder must not exceed that reported in the Local Clock Granularity field of the PTM Capability register.

6.x.2.3. PTM Time Source Role -- Rules Specific to Switches

In addition to the requirements listed above for the PTM Requester and PTM Responder Roles, Switches must follow these requirements:

- ❑ When the Upstream Port is associated with a multi-Function device, only a single Function associated with that Upstream Port is permitted to implement the PTM Extended Capability structure. For Switches, all PTM functionality associated with the Switch must be controlled through that structure. It is not required that the Function implementing the PTM Extended Capability structure be the Switch Upstream Port Function.
- ❑ The PTM Extended Capability structure for a Switch must indicate support for both the PTM Requester and PTM Responder roles.
- ❑ The PTM Extended Capability in the Upstream Port controls all Switches in that Upstream Port.
- ❑ A Switch is permitted to act as a PTM Root, or to issue PTM Requests on its Upstream Port to obtain the PTM Master Time for use in fulfilling PTM Requests received at its Downstream Ports. In the latter case the Switch must account for any internal delays within the Switch.
- ❑ A Switch is permitted to maintain a local PTM context for use in fulfilling PTM Requests received on its Downstream Ports.
- ❑ A Switch which is not acting as a PTM Root must invalidate its local context no more than 10 ms from the last PTM dialog on its Upstream Port. The Switch must then refresh its local PTM context prior to issuing further PTM ResponseD Messages on its Downstream Ports. This requirement for periodic refreshes is optional if it is guaranteed by implementation-specific means that the Switch's local clock is phase locked with PTM Master Time.

- Any Switch implementing a local clock for the purpose of maintaining a local PTM context must report the granularity of this clock as defined in the PTM Capabilities structure (section 7.x.3).



IMPLEMENTATION NOTE

PTM Timestamp Capture Mechanisms

PTM uses services from both the Data Link and Transaction Layers. Accuracy requires that time measurements be taken as close to the Physical Layer as possible. Conversely, the messaging protocol itself properly belongs to the Transaction Layer. The PTM message protocol applies to a single Link, where the Upstream Port is the requestor and the Downstream Port is the responder.

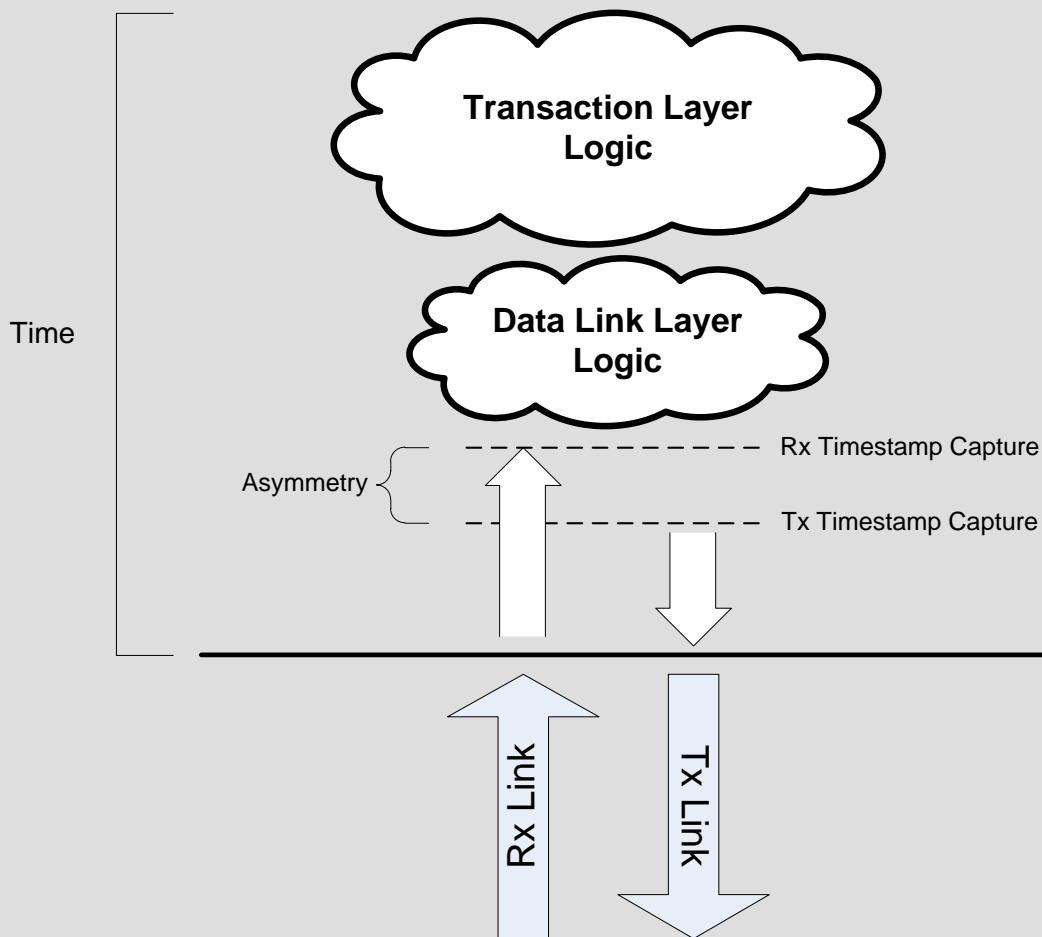


Figure 6-e: PTM Timestamp Capture Example

Figure 6-e illustrates how to select suitable timestamp capture points. For some implementations, the logic within the Transaction Layer and Data Link Layers is non-deterministic. Implementation details and current conditions have considerable impact on exactly when a particular packet may encounter any particular processing step. This makes it effectively impossible to capture any timestamp that accurately records the time of a particular physical event if timestamps are captured in the higher layers.

Add section 7.x as follows:

7.x. Precision Time Management (PTM) Capability

The Precision Time Measurement (PTM) Capability is an optional Extended Capability for discovering and controlling the distribution of a PTM Hierarchy. For Root Complexes, this Capability is required in any Root Port, Root Complex Integrated Endpoint, or RCRB that supports PTM. For Endpoints and Switch Upstream Ports that support PTM, this Capability is required in exactly one Function of the Upstream Port and that Capability controls the PTM behavior of all PTM capable Functions associated with that Upstream Port. For Switch Downstream Ports, PTM behavior is controlled by the same PTM Capability that controls the associated Switch Upstream Port. The PTM Capability is not permitted in Bridges, Switch Downstream Ports, and Root Complex Event Collectors.

For Switches, a single instance of this Capability controls behavior for the entire Switch. If the Upstream Port of the Switch is associated with an MFD, it is not required that the controlling Function be the Function corresponding to the Switch Upstream Port. For a given Switch, if this Capability is present, all Downstream Ports of the Switch must implement the requirements defined in Section 6.x.2.2.

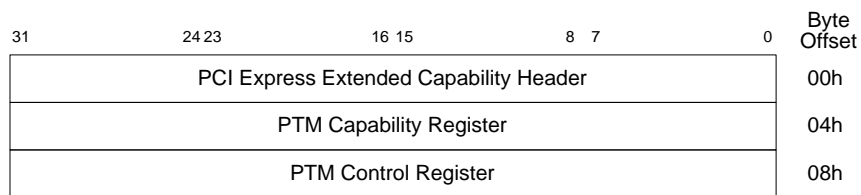


Figure 7-a: PTM Extended Capability Structure

7.x.1. PTM Extended Capability Header (Offset 00h)

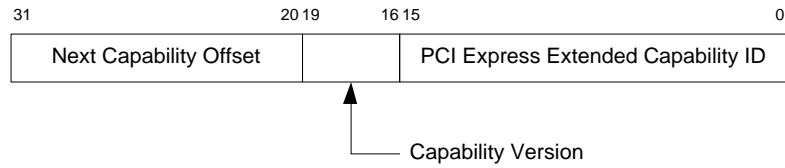


Figure 7-b: PTM Extended Capability Header

Table 7-x: PTM Extended Capability Header

<u>Bit Location</u>	<u>Register Description</u>	<u>Attributes</u>
<u>15:0</u>	<p>PCI Express Extended Capability ID – This field is a PCI-SIG defined ID number that indicates the nature and format of the Extended Capability.</p> <p>PCI Express Extended Capability ID for the Precision Time Measurement Capability is 001Fh.</p>	<u>HwInit</u>
<u>19:16</u>	<p>Capability Version – This field is a PCI-SIG defined version number that indicates the version of the Capability structure present.</p> <p>Must be 1h for this version of the specification.</p>	<u>HwInit</u>
<u>31:20</u>	<p>Next Capability Offset – This field contains the offset to the next PCI Express Extended Capability structure or 000h if no other items exist in the linked list of Capabilities.</p>	<u>HwInit</u>

7.x.2. PTM Capability Register (Offset 04h)

This register describes a Function's support for Precision Time Measurement. Not all fields within this register apply to all Functions capable of implementing PTM.

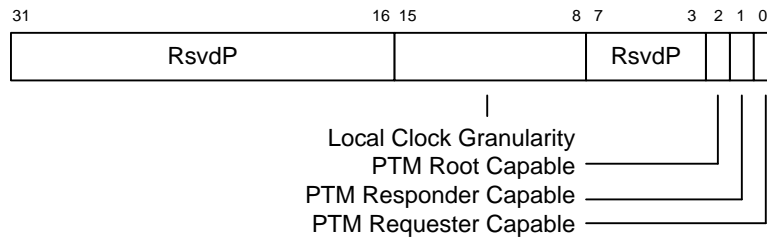


Figure 7-c: PTM Capability Register

Table 7-y: PTM Capability Register

<u>Bit Location</u>	<u>Register Description</u>	<u>Attributes</u>
<u>0</u>	<u>PTM Requester Capable</u> – Endpoints and Root Complex Integrated Endpoints are permitted to, and Switches supporting PTM must, set this bit to 1b to indicate they implement the PTM Requester role (see Section 6.x.2.1).	<u>HwInit</u>
<u>1</u>	<u>PTM Responder Capable</u> – Root Ports and RCRB's are permitted to, and Switches supporting PTM must, set this bit to 1b to indicate they implement the PTM Responder role (see Section 6.x.2.2). If PTM Root Capable is Set, this bit must be Set to 1b.	<u>HwInit</u>
<u>2</u>	<u>PTM Root Capable</u> – Root Ports, RCRB's and Switches are permitted to set this bit to 1b, if they implement a PTM Time Source Role and are capable of serving as the PTM Root.	<u>HwInit</u>
<u>7:3</u>	<u>Reserved</u>	<u>RsvdP</u>
<u>15:8</u>	<u>Local Clock Granularity</u> – <u>0000 0000b</u> – Time Source does not implement a local clock. It simply propagates timing information obtained from further Upstream in the PTM Hierarchy when responding to PTM Request messages. <u>0000 0001b</u> – <u>1111 1110b</u> : Indicates the period of this Time Source's local clock in ns. <u>1111 1111b</u> : Indicates the period of this Time Source's local clock is greater than 254 ns. If the PTM Root Select bit is Set, this local clock is used to provide PTM Master Time. Otherwise, the Time Source uses this local clock to locally track PTM Master Time received from further Upstream within a PTM Hierarchy. This field is reserved for Functions that do not implement the PTM Time Source role.	<u>HwInit/RsvdP</u>
<u>31:16</u>	<u>Reserved</u>	<u>RsvdP</u>

7.x.3. PTM Control Register (Offset 08h)

This register controls a Function's participation in the Precision Time Measurement mechanism. Not all fields within this register apply to all Functions capable of implementing PTM.

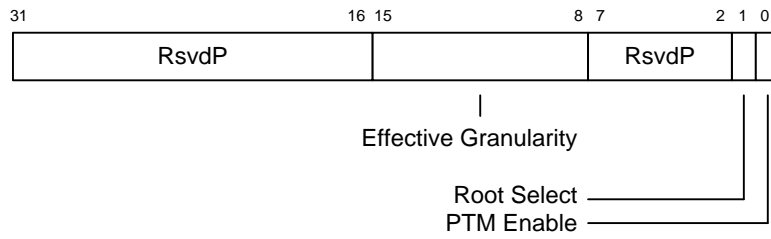


Figure 7-d: PTM Control Register

Table 7-z: PTM Control Register

<u>Bit Location</u>	<u>Register Description</u>	<u>Attributes</u>
<u>0</u>	<p><u>PTM Enable</u> – When Set, this Function is permitted to participate in the PTM mechanism according to its selected role(s) (see Section 6.x.2).</p> <p>Default value is 0b.</p>	<u>RW</u>
<u>1</u>	<p><u>Root Select</u> – When Set, if the PTM Enable bit is also Set, this Time Source is the PTM Root.</p> <p>Within each PTM Hierarchy, it is recommended that system software select only the furthest Upstream Time Source to be the PTM Root.</p> <p>Default value is 0b.If the value of the PTM Root Capable bit is 0b, this bit is permitted to be hardwired to 0b.</p>	<u>RW/RO</u>
<u>7:2</u>	<u>Reserved</u>	<u>RsvdP</u>
<u>15:8</u>	<p><u>Effective Granularity</u> – For Functions implementing the PTM Requester Role, this field provides information relating to the expected accuracy of the PTM clock, but does not otherwise affect the PTM mechanism.</p> <p>For Endpoints, system software must program this field to the value representing the maximum Local Clock Granularity reported by the PTM Root and all intervening PTM Time Sources.</p> <p>For Root Complex Integrated Endpoints, system software must set this field to the value reported in the Local Clock Granularity field by the associated PTM Time Source.</p> <p>Permitted values:</p> <p><u>0000 0000b</u> – Unknown PTM granularity – one or more Switches between this Function and the PTM Root reported a Local Clock Granularity value of <u>0000 0000b</u>.</p> <p><u>0000 0001b</u> – <u>1111 1110b</u> – Indicates the effective PTM granularity in ns.</p> <p><u>1111 1111b</u>: Indicates the effective PTM granularity is greater than 254 ns.</p> <p>Default value is <u>00000b</u>. If PTM Requester Capable is Clear, this field is permitted to be hardwired to <u>0000 0000b</u>.</p>	<u>RW/RO</u>
<u>31:16</u>	<u>Reserved</u>	<u>RsvdP</u>

Modify Table F-1 as follows:

Table F-1: Message Code Usage

Message Code	Routing R[2:0]	Type	Description
...
0101 0000	100	MsgD	Set_Slot_Power_Limit, see Section 2.2.8.5
<u>0101 0010</u>	<u>100</u>	<u>Msg</u>	<u>PTM Request, see Section 2.2.8.x</u>
<u>0101 0011</u>	<u>100</u>	<u>Msg/MsgD</u>	<u>PTM Response/PTM ResponseD, see Section 2.2.8.x</u>

In the PCI Code and ID Assignment Specification, Modify Table 3-1 as follows:

Table 3-1: Extended Capability IDs

ID	Extended Capability
...	...
<u>001Fh</u>	<u>Precision Time Measurement (PTM)</u>