PCI-SIG ENGINEERING CHANGE NOTICE

TITLE:	Extension Devices
DATE:	Updated 6 October 2014
AFFECTED DOCUMENT:	PCI Express Base Specification, Revision 3.0/3.1, Multi-Root I/O Virtualization and Sharing Specification, Revision 1.0. This ECN comprehends the following PCI Express Base Specification, Revision 3.0 ECNs: Separate Refclk Independent SSC (SRIS) Architecture, 8.0 GT/s Receiver Impedance, L1 PM Substates with CLKREQ
SPONSOR:	Intel, IDT

Part I

1. Summary of the Functional Changes

Provide specification for Physical Layer protocol aware Retimers for PCI Express 3.0/3.1.

2. Benefits as a Result of the Changes

Provides a consistent specification for a Physical Layer protocol aware Retimer that interoperates with all existing PCI Express 3.0/3.1 compliant devices.

3. Assessment of the Impact

There are no required changes for existing PCI Express 3.0/3.1 devices. New PCI Express 3.0/3.1 devices can optionally implement features to detect Retimer presence, report presence through configuration space and/or optionally adjust other parameters based on Retimer presence.

4. Analysis of the Software Implications

None.

5. Analysis of the C&I Test Implications

A new C&I program could be implemented to test standalone Retimer products that are implemented to this specification. No impact to testing of existing products that contain a Retimer.

Part II

Detailed Description of the change

Add the following entries to Terms and Abbreviations:

Downstream Path	The flow of data through a Retimer from the Upstream Pseudo Port Receiver to the Downstream Pseudo Port Transmitter.
…	
Extension Device	A Component whose function is to extend the physical length of a Link.
…	
Path	The flow of data through a Retimer, in either the Upstream Path or the Downstream Path.
…	
Pseudo Port	1. Logically, an interface between a Retimer and a PCI Express Link Segment. 2. Physically, a group of Transmitters and Receivers located on the same Retimer chip that define a Link Segment.
…	
Re-driver	A non-protocol aware, software transparent, analog only, Extension Device. Generally a device that does not contain a clock and data recovery (CDR), uses only a Continuous Time Linear Equalizer (CTLE) in the Receiver, and uses a fixed De-emphasis in the Transmitter.
…	
Repeater	A Retimer or a Re-driver.
…	
Retimer	A Physical Layer protocol aware, software transparent, Extension Device that forms two separate electrical Link Segments.
…	
Link Segment	The collection of a Port and a Pseudo Port or two Pseudo Ports and their interconnecting Lanes. A *Link Segment* is a dual simplex communications path between a Component and a Retimer or between two Retimers (two Pseudo Ports).
…	
Upstream Path	The flow of data through a Retimer from the Downstream Pseudo Port Receiver to the Upstream Pseudo Port Transmitter.

Modify Table 4-6

Table 4-6: TS2 Ordered Set

Symbol Number	Description
…	
5	Training Control Bit 0 – Hot Reset Bit 0 = 0b, De-assert Bit 0 = 1b, Assert Bit 1 – Disable Link Bit 1 = 0b, De-assert Bit 1 = 1b, Assert Bit 2 – Loopback Bit 2 = 0b, De-assert Bit 2 = 1b, Assert Bit 3 – Disable Scrambling in 2.5 GT/s and 5.0 GT/s data rates; Reserved in other data rates Bit 3 = 0b, De-assert Bit 3 = 1b, Assert Bit 4 – Retimer Present in 2.5 GT/s data rate. Reserved in other data rates. Bit 4 = 0b, No Retimers present Bit 4 = 1b, One or more Retimers present Bit ~~4:7~~ 5:7 – Reserved
…	

Modify Section 4.2.6.3.5.1:

…

- The next state is Configuration.Idle immediately after all Lanes that are transmitting TS2 Ordered Sets receive eight consecutive TS2 Ordered Sets with matching Lane and Link numbers (non-PAD) and identical data rate identifiers (including identical Link Upconfigure Capability (Symbol 4 bit 6)), and 16 TS2 Ordered Sets are sent after receiving one TS2 Ordered Set. Implementations with the Retimer Presence Detect Supported bit of the Link Capabilities 2 register set to 1b must also receive the eight consecutive TS2 Ordered Sets with identical Retimer Present (Symbol 5 bit 4) when the data rate is 2.5 GT/s.
  - If the data rate of operation is 2.5 GT/s:
    - If the Retimer Presence Detect Supported bit of the Link Capabilities 2 register is set to 1b and any Configured Lane received the Retimer Present bit set to 1b in the eight consecutively received TS2 Ordered Sets then the Retimer Presence Detected bit must be set to 1b in the Link Status 2 Register otherwise the Retimer Presence Detected bit must be set to 0b in the Link Status 2 Register.

…

Modify Section 4.2.7.1:

4.2.7.1 SKP Ordered Set for 8b/10b Encoding

When using 8b/10b encoding, a transmitted SKP Ordered Set is a COM Symbol followed by three SKP Symbols, except as is allowed for a Loopback Slave in the Loopback.Active LTSSM state. A received SKP Ordered Set is a COM Symbol followed by one to five SKP Symbols. See Section 4.4.6.7 for Retimer rules on SKP Ordered Set modification.

Modify Section 4.2.7.2:

4.2.7.2 SKP Ordered Set for 128b/130b Encoding

Table 4-15 describes the layout of the SKP Ordered Set when using 128b/130b encoding. A transmitted SKP Ordered Set is 16 Symbols. Four SKP Symbols can be added and removed by a Port. Hence, a received SKP Ordered Set can be 8, 12, 16, 20, or 24 Symbols. The SKP_END Symbol indicates the last four Symbols of SKP Ordered Set so that Receivers can identify the location of the next Block Sync Header in the bit stream. The three Symbols following the SKP_END Symbol contain different information depending on the LTSSM state. See Section 4.4.6.7 for Retimer rules on SKP Ordered Set modification.

Add Section 4.4:

4.4 Retimers

This Section defines the requirements for Retimers that are Physical Layer protocol aware and that interoperate with any pair of Components with any compliant channel on each side of the Retimer. This Retimer definition does not support M-PCIe, See Chapter 8 for M-PCIe details. An important capability of a Physical Layer protocol aware Retimer is to execute the Phase 2/3 of the equalization procedure in each direction. A maximum of two Retimers are permitted between an Upstream and a Downstream Component.

The two Retimer limit is based on multiple considerations, most notably limits on modifying SKP Ordered Sets and limits on the time spent in Phase 2/3 of the equalization procedure. To ensure interoperability, platform designers must ensure that the two Retimer limit is honored for all PCI Express Links, including those involving form factors as well as those involving active cables. Form factor specifications may define additional Retimer rules that must be honored for their form factors. Assessing interoperability with any Extension Device not based on the Retimer definition in this section is outside the scope of this specification.

Many architectures of Extension Devices are possible, i.e., analog only Repeater, protocol unaware Retimer, etc. This specification describes a Physical Layer protocol aware Retimer. It may be possible to use other types of Extension Devices in closed systems if proper analysis is done for the specific channel, Extension Device, and end-device pair – but a specific method for carrying out this analysis is outside the scope of this specification.

Retimers have two Pseudo Ports, one facing Upstream, and the other facing Downstream. The Transmitter of each Pseudo Port must derive its clock from a 100 MHz reference clock. The reference clock(s) must meet the requirements of Section 4.3.7 and Section 4.3.8. A Retimer supports one or more reference clocking architectures as defined in Section 4.3 Electrical Sub-block.

In most operations Retimers simply forward received Ordered Sets, DLLPs, TLPs, Logical Idle, and Electrical Idle. Retimers are completely transparent to the Data Link Layer and Transaction Layer. System software shall not enable L0s on any Link where a Retimer is present. Support of beacon by Retimers is optional and beyond the scope of this specification.

When using 128b/130b encoding the Retimer executes the protocol so that each Link Segment undergoes independent Link equalization as described in Section 4.4.7.2.

The Pseudo Port orientation (Upstream or Downstream), is determined dynamically, while the Link partners are in Configuration. Both crosslink and regular Links are supported.

4.4.1 Retimer Requirements

The following is a high level summary of Retimer requirements:

- Retimers are required to comply with all the electrical specification described in Section 4.3 Electrical Sub-block. Retimers' receivers must operate at 8.0 GT/s with an impedance that meets the range defined by the Z_{RX-DC} parameter for 2.5 GT/s.
- Forwarded Symbols must always be de-skewed when more than one Lane is forwarding Symbols (including upconfigure cases).

- ❑ Determine Port orientation dynamically.
- ❑ Perform Lane polarity inversion (if needed).
- ❑ Execute the Link equalization procedure for Phase 2 and Phase 3, when using 128b/130b encoding, on each Link Segment.
- ❑ Interoperate with de-emphasis negotiation at 5.0 GT/s, on each Link Segment.
- ❑ Interoperate with Link Upconfigure
- ❑ Pass loopback data between the loopback master and loopback slave.
    - Optionally execute Slave Loopback on one Pseudo Port.
- ❑ Generate the Compliance Pattern on each Pseudo Port.
    - Load board method (i.e., time out in Polling.Active).
- ❑ Forward Modified Compliance Pattern when the Link enters Polling.Compliance via Compliance Receive bit in TS1 Ordered Sets.
- ❑ Forward Compliance or Modified Compliance Patterns when Components enter Polling.Compliance via the Enter Compliance bit in the Link Control 2 register is set to 1b in both the Upstream Component and the Downstream Component and Retimer Enter Compliance is set to 1b (accessed in an implementation specific manner) in the Retimer.
- ❑ Adjust the data rate of operation in concert with the Upstream and Downstream Components of the Link.
- ❑ Adjust the Link width in concert with the Upstream and Downstream Components of the Link.
- ❑ Capture Lane numbers during Configuration.
    - Lane numbers are required when using 128b/130b encoding for the scrambling seed.
- ❑ Dynamically adjust Retimer Receiver impedance to match end Component Receiver impedance.
- ❑ Infer entering Electrical Idle at all data rates.
- ❑ Modify certain fields of Ordered Sets while forwarding.
- ❑ Perform clock compensation via addition or removal of SKP Symbols.
- ❑ Support L1.
    - Optionally Support L1 PM Substates.

4.4.2 Supported Topologies

Figure 4-a Supported Topologies shows the topologies supported by Retimers defined in this specification. There may be one or two Retimers between the Upstream and Downstream Components on a Link. Each Retimer has two Pseudo Ports, which determine their Downstream/Upstream orientation dynamically. Each Retimer has an Upstream Path and a Downstream Path. Both Pseudo Ports must always operate at the same data rate, when in Forwarding mode. Thus each Path will also be at the same data rate. A Retimer is permitted

to support any width option defined by this specification as its maximum width. The behavior of the Retimer in each high level operating mode is:

- Forwarding mode:
  - Symbols, Electrical Idle, and exit from Electrical Idle; are forwarded on each Upstream and Downstream Path.
- Execution mode:
  - The Upstream Pseudo Port acts as an Upstream Port of a Component. The Downstream Pseudo Port acts as a Downstream Port of a Component. This mode is used in the following cases:
    - Polling.Compliance.
    - Phase 2 and Phase 3 of the Link equalization procedure.
    - Optionally Slave Loopback.

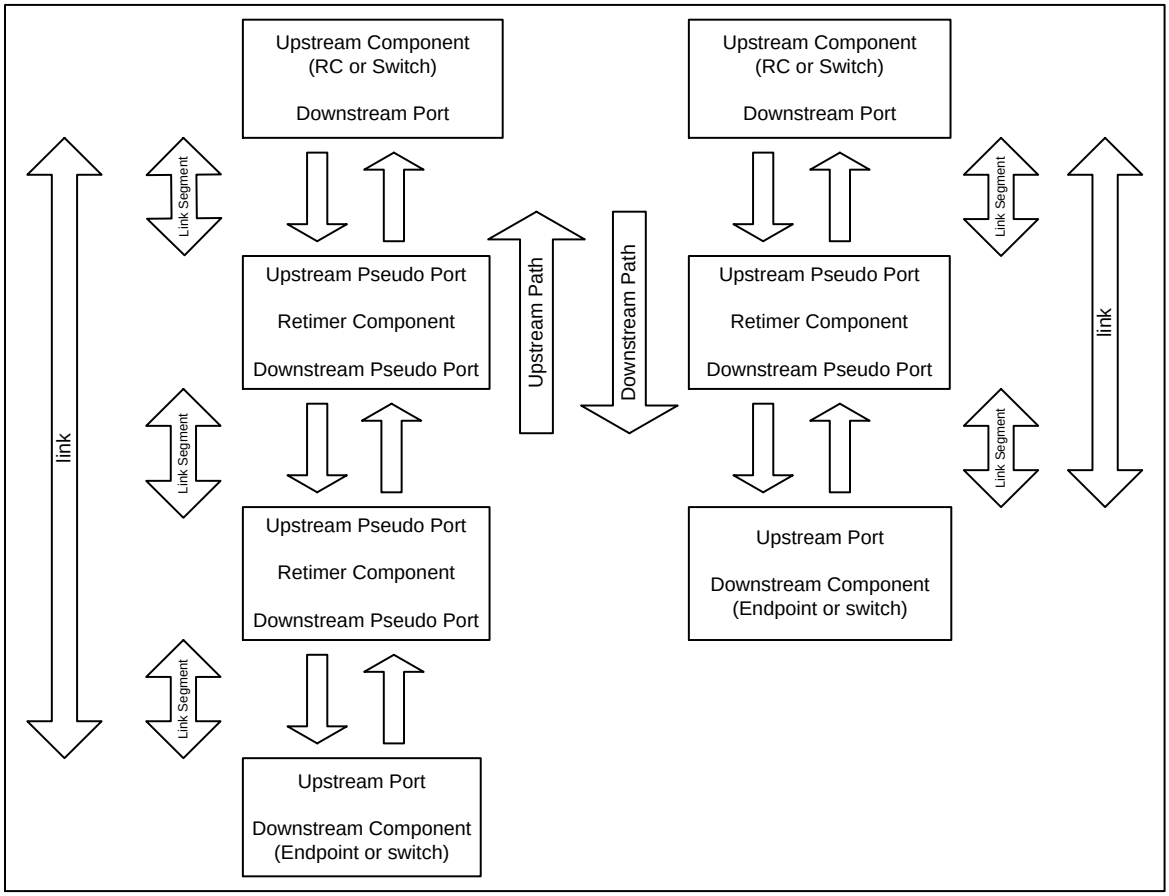

Figure 4-a Supported Topologies

4.4.3 Variables

The following variables are set to the following specified values following a Fundamental Reset or whenever the Retimer receives Link and Lane number equal to PAD on two

consecutive TS2 Ordered Sets on all Lanes that are receiving TS2 Ordered Sets on both Upstream and Downstream Pseudo Ports within a 1 µs time window from the last Symbol of the second TS2 Ordered Set on the first Lane to the last Symbol of the second TS2 Ordered Set on the last Lane.

- ❑ RT_port_orientation = undefined
- ❑ RT_captured_lane_number = PAD
- ❑ RT_captured_link_number = PAD
- ❑ RT_G3_EQ_complete = 0b
- ❑ RT_LinkUp = 0b
- ❑ RT_next_data_rate = 2.5 GT/s
- ❑ RT_error_data_rate = 2.5 GT/s

4.4.4 Receiver Impedance Propagation Rules

The Retimer Transmitters and Receivers shall meet the requirements in Section 4.2.4.8.1 while Fundamental Reset is asserted. When Fundamental Reset is de-asserted the Retimer is permitted to take up to 20 ms to begin active determination of its Receiver impedance. During this interval the Receiver impedance remains as required during Fundamental Reset. Once this interval has expired Receiver impedance on Retimer Lanes is determined as follows:

- ❑ Within 1.0 ms of the Upstream or Downstream Component's Receiver meeting Z_{RX-DC}, the low impedance is back propagated, i.e., the Retimer's Receiver shall meet Z_{RX-DC} on the corresponding Lane on the other Pseudo Port. Each Lane operates independently and this requirement applies at all times.
- ❑ The Retimer must keep its Transmitter in Electrical Idle until the Z_{RX-DC} state has been detected. This applies on an individual Lane basis.

4.4.5 Switching Between Modes

The Retimer operates in two basic modes, Forwarding mode or Execution mode. When switching between these modes the switch must occur on an Ordered Set boundary for all Lanes of the Transmitter at the same time. No other Symbols shall be between the last Ordered Set transmitted in the current mode and the first Symbol transmitted in the new mode.

When using 128b/130b the Transmitter must maintain the correct scrambling seed and LFSR value when switching between modes.

When switching to and from the Execution Link Equalization mode the Retimer must ensure a Transmitter does not send two SKP Ordered Sets in a row, and that the maximum allowed interval is not exceeded between SKP Ordered Sets, see Section 4.2.7.3.

4.4.6 Forwarding Rules

These rules apply when the Retimer is in Forwarding mode. The Retimer is in Forwarding mode after the deassertion of Fundamental Reset.

- ❑ If the Retimer's Receiver detects an exit from Electrical Idle on a Lane the Retimer must enter Forwarding mode and forward the Symbols on that Lane to the opposite Pseudo Port as described in Section 4.4.6.3.

- ❑ The Retimer must continue to forward the received Symbols on a given Lane until it enters Execution mode or until an EIOS is detected, or until Electrical Idle is inferred on that Lane. This requirement applies even if the Receiver loses Symbol lock or Block Alignment. See Section 4.4.6.5 for rules regarding Electrical Idle entry.

- ❑ A Retimer shall forward all Symbols unchanged, except as described in Section 4.4.6.9 and 4.4.6.7.

- ❑ When operating at 2.5 GT/s data rate, if any Lane of a Pseudo Port receives TS1 Ordered Sets with Link and Lane numbers set to PAD for 5 ms or longer, and the other Pseudo Port does not detect an exit from Electrical Idle on any Lane in that same window, and either of the following occurs:

  - An EIOS is received on any Lane that was receiving TS1 Ordered Sets (this is interpreted as the Component attached to the Receiver going into Electrical Idle prior to a data rate change for a Compliance Pattern above 2.5 GT/s).

  - Compliance Pattern at 2.5 GT/s is received on any Lane that was receiving TS1 Ordered Sets.

  Then the Retimer enters the Execution mode CompLoadBoard state, and follows Section 4.4.7.1.

  If TS1 Ordered Sets at 2.5 GT/s, with Lane and Link number equal to PAD are received on any Lane of one Pseudo Port for 12 ms (+/- 1.0 ms) and the other Pseudo Port does not detect an exit from Electrical Idle on any Lane in the same 12 ms window then the Retimer must stop forwarding the TS1 Ordered Sets on the Port that was forwarding the Ordered Sets and performs Receiver detection. If the Receiver detection detects the same Lanes that previously were detected then the Retimer resumes forwarding TS1 Ordered sets, else the Transmitter sends 32 K28.7 Symbols (EIE) on all lanes of the Pseudo Port that are receiving TS1 Orderd Sets and then sets all receiver terminations to high impedance and returns to the conditions described in Section 4.4.4 following Fundamental Reset.

- ❑ If any Lane on the Upstream Pseudo Port receives two consecutive TS1 Ordered Sets with the EC field equal to 10b, when using 128b/130b encoding, then the Retimer enters Execution mode Equalization, and follows Section 4.4.7.2.

- ❑ If the Retimer is configured to support Execution mode Slave Loopback and if any Lane on either Pseudo Port receives two consecutive TS1 Ordered Sets or two consecutive TS2 Ordered Sets with the Loopback bit set to 1b then the Retimer enters Execution mode Slave Loopback, and follows Section 4.4.7.3.

4.4.6.1 Forwarding Type Rules

A Retimer must determine what type of Symbols it is forwarding. The rules for inferring Electrical Idle are a function of the type of Symbols the Retimer is forwarding. If a Path forwards two consecutive TS1 Ordered Sets or two consecutive TS2 Ordered Sets, on any Lane, then the Path is forwarding training sets. If a Path forwards eight consecutive Symbol Times of Idle data on all Lanes that are forwarding Symbols then the Path is forwarding non-training sets.

4.4.6.2 Orientation and Lane Numbers Rules

The Retimer must determine the Port orientation, Lane assignment, and Lane polarity dynamically as the Link trains.

- When RT_LinkUp=0, the first Pseudo Port to receive two consecutive TS1 Ordered Sets with a non-PAD Lane number on any Lane, has its RT_port_orientation variable set to Upstream Port, and the other Pseudo Port has its RT_port_orientation variable set to Downstream Port.

- The Retimer plays no active part of Lane number determination. The Retimer must capture the Lane numbers with the RT_captured_lane_number variable at the end of the Configuration state, between the Link Components. This applies on the first time through Configuration, i.e., when RT_LinkUp is set to 0b. Subsequent trips through Configuration during Link width configure must not change the Lane numbers. Lane numbers are required for the scrambling seed when using 128b/130b. Link numbers are required in some cases when the Retimer is in Execution mode. Link numbers and Lane numbers are captured with the RT_captured_lane_number, and RT_captured_link_number variables whenever the first two consecutive TS2 Ordered Sets that contain non-PAD Lane and non-PAD Link numbers are received after RT_LinkUp variable is set to 0b. A Retimer must function normally if Lane reversal occurs. When the Retimer has captured the Lane numbers and Link numbers the variable RT_LinkUp is set to 1b. In addition if the Disable Scrambling bit in the TS2 Ordered Sets is set to 1b, in either case above, then the Retimer determines that scrambling is disabled when using 8b/10b encoding.

- Lane polarity is determined any time the Lane exits Electrical Idle, and achieves Symbol lock at 2.5 GT/s as described in Section 4.2.4.4:

  - If polarity inversion is determined the Receiver must invert the received data. The Transmitter must never invert the transmitted data.

4.4.6.3 Electrical Idle Exit Rules

At data rates other than 2.5 GT/s, EIEOS are sent within the training sets to ensure that the analog circuit detects an exit from Electrical Idle. Receiving an EIEOS is required when using 128b/130b encoding to achieve Block Alignment. When the Retimer starts forwarding data after detecting an Electrical Idle exit, the Retimer starts transmitting on a training set boundary. The first training sets it forwards must be an EIEOS, when operating at data rates higher than 2.5 GT/s. The first EIEOS sent will be in place of the TS1 or TS2 Ordered Set that it would otherwise forward.

If a Lane detects an exit from Electrical Idle then the Lane must start forwarding when all of the following are true:

- ❑ Data rate is determined, see Section 4.4.6.4, current data rate is changed to RT_next_data_rate if required.
- ❑ Lane polarity is determined, see Section 4.4.6.2.
- ❑ Two consecutive TS1 Ordered Sets or two consecutive TS2 Ordered Sets are received.
- ❑ Two consecutive TS1 Ordered Sets or two consecutive TS2 Ordered Sets are received on all Lanes that detected an exit from Electrical Idle or the max Retimer Exit Latency has occurred, see Table 4-a.
- ❑ Lane De-skew is achieved on all Lanes that received two consecutive TS1 or two consecutive TS2 Ordered Sets.

  All Ordered Sets used to establish forwarding must be discarded.

Otherwise after a 3.0 ms timeout, the pattern described below is transmitted on all Lanes that detected an exit from Electrical Idle.

- ❑ When using 128b/130b encoding:
  - One EIEOS
  - 32 Data Blocks, each with a payload of 16 Idle data Symbols (00h), scrambled, for Symbols 0 to 13.
  - Symbol 14 and 15 of each Data Block either contain Idle data Symbols (00h), scrambled, or DC Balance, determined by applying the same rules in 4.2.4.1 to these Data Blocks.
- ❑ When using 8b/10b encoding:
  - The Modified Compliance Pattern with the error status Symbol set to 00h.

This Path now is forwarding Electrical Idle Exit. In this state Electrical Idle is inferred by the absence of Electrical Idle Exit, See Table 4-b.

IMPLEMENTATION NOTE

Electrical Idle Exit

Forwarding Electrical Idle Exit occurs in error cases where a Retimer is unable to decode training sets. Upstream and Downstream Components use Electrical Idle Exit (without decoding any Symbols) during Polling.Compliance, and Recovery.Speed. If the Retimer does not forward Electrical Idle Exit then the Upstream and Downstream Components will misbehave in certain conditions. For example, this may occur after a speed change to a higher data rate. In this event forwarding Electrical Idle Exit is required to keep the Upstream and Downstream Components in lock step at Recovery.Speed, so that the data rate will return to the previous data rate, rather than a Link Down condition from a time out to Detect.

When a Retimer detects an exit from Electrical Idle and starts forwarding data, the time this takes is called the Retimer Exit Latency, and allows for such things as data rate change (if required), clock and data recovery, Symbol lock, Block Alignment, Lane-to-Lane de-skew, Receiver tuning, etc. The maximum Retimer Exit Latency is specified below for several conditions:

- The data rate before and after Electrical Idle and Electrical Idle exit detect does not change.
- Data rate change to a data rate that uses 8b/10b encoding.
- Data rate change to a data rate that uses 128b/130b encoding for the first time.
- Data rate change to a data rate that uses 128b/130b encoding not for the first time.
- How long both transmitters have been in Electrical Idle when a data rate change occurs.

Retimers are permitted to change their data rate while in Electrical Idle, and it is recommended that Retimers start the data rate change while in Electrical Idle to minimize Retimer Exit latency.

Table 4-a: Maximum Retimer Exit Latency

Condition	**Link in EI** **For X us, where** **X < 500 us**	**Link in EI for** **For X ≥ 500 us**
No data rate change	4 us	4 us
Any data rate change to 8b/10b encoding data rate	500 – X us	4 us
First data rate change to 128b/130b encoding date rate	1.5 -X ms	1 ms
Subsequent data rate change to 128b/130b encoding date rate	500 –X us	4 us

4.4.6.4 Data Rate Change and Determination Rules

The data rate of the Retimer is set to 2.5 GT/s after deassertion of Fundamental Reset.

Both Pseudo Ports of the Retimer must operate at the same data rate. If a Pseudo Port places its Transmitter in Electrical Idle, then the Symbols that it has just completed transmitting determine the variables RT_next_data_rate and RT_error_data_rate. Only when both Pseudo Ports have all Lanes in Electrical Idle shall the Retimer change the data rate. If both Pseudo Ports do not make the same determination of these variables then both variables must be set to 2.5 GT/s.

- If both Pseudo Ports were forwarding non-training sequences, then the RT_next_data_rate must be set to the current data rate. The RT_error_data_rate must be set to 2.5 GT/s.

- If both Pseudo Ports were forwarding TS2 Ordered Sets and the highest common data rate is not equal to the current data rate, then RT_next_data_rate must be set to the highest common data rate.

- If at least one Pseudo Port was forwarding Electrical Idle exit, then the RT_next_data_rate must be set to the RT_error_data_rate. The RT_error_data_rate is set to 2.5 GT/s.

4.4.6.5 Electrical Idle Entry Rules

The Rules for Electrical Idle entry in Forwarding mode are a function of whether the Retimer is forwarding training sets or non-training sets. The determination of this is described in Section 4.4.6.1.

Before a Transmitter enters Electrical Idle, it must always send the Electrical Idle Ordered Set (EIOS), unless otherwise specified.

If the Retimer is forwarding training sets then:

- If an EIOS is received on a Lane, then the EIOS is forwarded on that Lane and only that Lane places its Transmitter in Electrical Idle. When operating at 5.0 GT/s, two EIOS are transmitted prior to entering Electrical Idle.

- If Electrical Idle is inferred on a Lane, then that Lane places its Transmitter in Electrical Idle, after EIOS is transmitted on that Lane. When operating at 5.0 GT/s, two EIOS are transmitted prior to entering Electrical Idle.

Else if the Retimer is forwarding non-training sets then:

- If an EIOS is received on any Lane, then the EIOS is forwarded on all Lanes that are currently forwarding Symbols and all Lanes place their Transmitters in Electrical Idle. When operating at 5.0 GT/s, two EIOS are transmitted prior to entering Electrical Idle.

- If Electrical Idle is inferred on a Lane, then that Lane places its Transmitter in Electrical Idle, and EIOS is not transmitted on that Lane.

The Retimer is required to infer Electrical Idle. The criteria for a Retimer inferring Electrical Idle are described in Table 4-b.

Table 4-b: Inferring Electrical Idle

State	2.5 GT/s	5.0 GT/s	8.0 GT/s
Forwarding: Non Training Sequence	Absence of a SKP Ordered Set in a 128 µs window	Absence of a SKP Ordered Set in a 128 µs window	Absence of a SKP Ordered Set in a 128 µs window
Forwarding: Training Sequence	Absence of a TS1 or TS2 Ordered Set in a 1280 UI interval	Absence of a TS1 or TS2 Ordered Set in a 1280 UI interval	Absence of a TS1 or TS2 Ordered Set in a 4680 UI interval
Forwarding: Electrical Idle Exit Executing: Force Timeout	Absence of an exit from Electrical Idle in a 2000 UI interval	Absence of an exit from Electrical Idle in a 16000 UI interval	Absence of an exit from Electrical Idle in a 16000 UI interval
Forwarding: Loopback Executing: Loopback Slave	Absence of an exit from Electrical Idle in a 128 µs window	N/A	N/A

4.4.6.6 Transmitter Settings Determination Rules

When a data rate change to 8.0 GT/s occurs the Retimer transmitter settings are determined as follows:

- ❑ If the RT_G3_EQ_complete variable is set to 1b:
  - The Transmitter must use the coefficient settings agreed upon at the conclusion of the last equalization procedure.
- ❑ Else:
  - An Upstream Pseudo Port must use the preset values it registered from the received eight consecutive TS2 Ordered Sets in its Transmitter setting as soon as it starts transmitting at the 8.0 GT/s data rate and must ensure that it meets the preset definition in Section 4.3.
  - A Downstream Pseudo Port determines its Transmitter Settings in an implementation specific manner when it starts transmitting at 8.0 GT/s.

The RT_G3_EQ_complete variable is set to 1b when:

- ❑ Two consecutive TS1 Ordered Sets are received with EC = 01b at 8.0 GT/s.

The RT_G3_EQ_complete variable is set to 0b when any of the following occur:

- RT_LinkUp variable is set to 0b
- Eight consecutive EQ TS1 or eight consecutive EQ TS2 Ordered Sets, with identical values in Symbol 6, are received on any Lane of the Upstream Pseudo Port. The value in Symbol 6 is registered for later use at 8.0 GT/s for that Lane.

When a data rate change to 5.0 GT/s occurs the Retimer transmitter settings are determined as follows:

- The Upstream Pseudo Port must sets its Transmitters to either -3.5 dB or -6.0 dB, according to the Selectable De-emphasis bit (bit 6 of Symbol 4) received in eight consecutive TS2 Ordered Sets, in the most recent series of TS2 Ordered sets, received prior to entering Electrical Idle.
- The Downstream Pseudo Port sets its Transmitters to either -3.5 dB or -6.0 dB in an implementation specific manner.

4.4.6.7 Ordered Set Modification Rules

Ordered Sets are forwarded, and certain fields are modified according to the following rules:

- The Retimer shall not modify any fields except those specifically allowed/required for modification in this specification.
- LF: the Retimer shall overwrite the LF field in the TS1 Ordered Sets transmitted in both directions. The new value is determined in an implementation specific manner by the Retimer.
- FS: the Retimer shall overwrite the FS field in the TS1 Ordered Sets transmitted in both directions. The new value is determined in an implementation specific manner by the Retimer.
- Transmitter Preset: the Retimer is permitted to overwrite the Transmitter Preset in the EQ TS1 Ordered Set or EQ TS2 Ordered Set in the Downstream Path. The new value is determined in an implementation specific manner by the Retimer.
- Receiver Preset Hint: the Retimer is permitted to overwrite the Receiver Preset Hint in the EQ TS1 Ordered Set or EQ TS2 Ordered Set in the Downstream Path. The new value is determined in an implementation specific manner by the Retimer.
- SKP Ordered Set: The Retimer is permitted to adjust the length of SKP Ordered Sets transmitted in both directions. The Retimer must perform the same adjustment on all Lanes. When operating with 8b/10b encoding, the Retimer is permitted to add or remove one SKP Symbol of a SKP Ordered Set. When operating with 128b/130b encoding, a Retimer is permitted to add or remove 4 SKP Symbols of a SKP Ordered Set.
- Selectable De-emphasis: the Retimer is permitted to overwrite the Selectable De-emphasis field in the TS1 or TS2 Ordered Set in both directions. The new value is determined in an implementation specific manner by the Retimer.

- The Data Rate Identifier: The Retimer must set the Data Rate Supported bits of the Data Rate Identifier Symbol consistent with the data rates advertised in the received Ordered Sets and its own Max Supported Data Rate. A Retimer must support all data rates below and including its advertised max data rate.
- DC Balance: When operating with 128b/130b encoding, the Retimer tracks the DC Balance of its Pseudo Port transmitters and transmits DC Balance Symbols as described in Section 4.2.4.1.
- Retimer Present: When operating at 2.5 GT/s, the Retimer must set the Retimer Present bit of all forwarded TS2 Ordered Sets to 1b.

4.4.6.8 DLLP, TLP, and Logical Idle Modification Rules

DLLPs, TLPs, and Logical Idle are forwarded with no modifications to any of the Symbols unless otherwise specified.

4.4.6.9 8b/10b Encoding Rules

The Retimer shall meet the requirements in Section 4.2.1.1.3 except as follows:

- When the Retimer is forwarding and an 8b/10b decode error or a disparity error is detected in the received data, the Symbol with an error is replaced with the D21.3 Symbol with incorrect disparity in the forwarded data.
- This clause in Section 4.2.1.1.3 does not apply: If a received Symbol is found in the column corresponding to the incorrect running disparity or if the Symbol does not correspond to either column, the Physical Layer must notify the Data Link Layer that the received Symbol is invalid. This is a Receiver Error, and is a reported error associated with the Port (see Section 6.2).

IMPLEMENTATION NOTE

Retimer Transmitter Disparity

The Retimer must modify certain fields of the TS1 and TS2 Ordered Sets (e.g., Receiver Preset Hint, Transmitter Preset), therefore the Retimer must recalculate the running disparity. Simply using the disparity of the received Symbol may lead to an error in the running disparity. For example some 8b/10b codes have 6 ones and 4 zeros for positive disparity, while other codes have 5 ones and 5 zeros.

4.4.6.10 8b/10b Scrambling Rules

A Retimer is required to determine if scrambling is disabled when using 8b/10b encoding as described in Section 4.4.6.2.

4.4.6.11 Hot Reset Rules

If any Lane of the Upstream Pseudo Port receives two consecutive TS1 Ordered Sets with the Hot Reset bit set to 1b and both the Disable Link and Loopback bits set to 0b, and then both Pseudo Ports either receive an EIOS or infer Electrical Idle on any Lane, that is receiving TS1 Ordered Sets, the Retimer does the following:

- ❑ Clears variable RT_LinkUp = 0b.
- ❑ Places its Transmitters in Electrical Idle on both Pseudo Ports.
- ❑ Set the RT_next_data_rate variable to 2.5 GT/s.
- ❑ Set the RT_error_data_rate variable to 2.5 GT/s.
- ❑ Waits for an exit from Electrical Idle on every Lane on both Pseudo Ports.

The Retimer does not perform Receiver detection on either Pseudo Port.

4.4.6.12 Disable Link Rules

If any Lane of the Upstream Pseudo Port receives two consecutive TS1 Ordered Sets with the Disable Link bit set to 1b and both the Hot Reset and Loopback bits set to 0b, and then both Pseudo Ports either receive an EIOS or infer Electrical Idle on any Lane, that is receiving TS1 Ordered Sets, the Retimer does the following:

- ❑ Clears variable RT_LinkUp = 0b.
- ❑ Places its Transmitters in Electrical Idle on both Pseudo Ports.
- ❑ Set the RT_next_data_rate variable to 2.5 GT/s.
- ❑ Set the RT_error_data_rate variable to 2.5 GT/s.
- ❑ Waits for an exit from Electrical Idle on every Lane on both Pseudo Ports.

The Retimer does not perform Receiver detection on either Pseudo Port.

4.4.6.13 Loopback

The Retimer follows these additional rules if any Lane receives two consecutive TS1 Ordered Sets with the Loopback bit equal to 1b and both the Hot Reset and Disable Link bits set to 0b and the ability to execute Slave Loopback is not configured in an implementation specific way. The purpose of these rules is to allow interoperation when a Retimer (or two Retimers) exist between a Loopback master and a Loopback slave.

- ❑ The Pseudo Port that received the TS1 Ordered Sets with the Loopback bit set to 1b acts as the Loopback Slave (the other Pseudo Port acts as Loopback Master). The Upstream Path is defined as the Pseudo Port that is the Loopback master to the Pseudo Port that is the Loopback slave. The other Path is the Downstream Path.
- ❑ Once established, if a Lane loses the ability to maintain Symbol Lock or Block alignment, then the Lane must continue to transmit Symbols while in this state.

- ❑ When using 8b/10b encoding and Symbol lock is lost, the Retimer must attempt to re-achieve Symbol Lock.

- ❑ When using 128b/130b encoding and Block Alignment is lost, the Retimer must attempt to re-achieve Block Alignment via SKP Ordered Sets.

- ❑ If Loopback was entered while the Link Components were in Configuration.Linkwidth.Start, then determine the highest common data rate of the data rates supported by the Link via the data rates received in two consecutive TS1 Ordered Sets or two consecutive TS2 Ordered Sets on any Lane, that was receiving TS1 or TS2 Ordered Sets, at the time the transition to Forwarding.Loopback occurred. If the current data rate is not the highest common data rate, then:

    - Wait for any Lane to receive EIOS, and then place the Transmitters in Electrical Idle for that Path.

    - When all Transmitters are in Electrical Idle, adjust the data rate as previously determined.

    - If the new data rate is 5.0 GT/s, then the Selectable De-emphasis is determined the same as way as described in Section 4.2.6.10.1.

    - If the new data rate uses 128b/130b encoding, then the Transmitter preset is determined the same as way as described in Section 4.2.6.10.1.

    - In the Downstream Path; wait for Electrical Idle exit to be detected on each Lane and then start forwarding when two consecutive TS1 Ordered Sets have been received, on a Lane by Lane basis. This is considered the first time to this data rate for the Retimer exit latency.

    - In the Upstream Path; if the Compliance Receive bit of the TS1 Ordered Sets that directed the slave to this state was not asserted, then wait for Electrical Idle exit to be detected on each Lane, and start forwarding when two consecutive TS1 Ordered Sets have been received, on a Lane by Lane basis. This is considered the first time to this data rate for the Retimer exit latency.

- ❑ In the Upstream Path; if the Compliance Receive bit of the TS1 Ordered Sets that directed the slave to this state was set to 1b, then wait for Electrical Idle exit to be detected on each Lane, and start forwarding immediately, on a Lane by Lane basis. This is considered the first time to this data rate for the Retimer exit latency.

- ❑ If four EIOS (one EIOS if the current data rate is 2.5 GT/s) are received on any Lane then:

    - Transmit eight EIOS on every Lane that is transmitting TS1 Ordered Sets on the Pseudo Port that did not receive the EIOS and place the Transmitters in Electrical Idle.

- ❑ When both Pseudo Ports have placed their Transmitters in Electrical Idle then:

    - Set the RT_next_data_rate variable to 2.5 GT/s.

- Set the RT_error_data_rate variable to 2.5 GT/s.
- The additional rules for Loopback no longer apply unless the rules for entering this Section are met again.

4.4.6.14 Compliance Receive Rules

The Retimer follows these additional rules if any Lane receives eight consecutive TS1 Ordered Sets (or their complement) with the Compliance Receive bit set to 1b and the Loopback bit set to 0b. The purpose of the following rules is to support Link operation with a Retimer when the Compliance Receive bit is set to 1b and the Loopback bit is set to 0b in TS1 Ordered Sets, transmitted by the Upstream or Downstream Component, while the Link is in Polling.Active.

- Pseudo Port A is defined as the first Pseudo Port that receives eight consecutive TS1 Ordered Sets (or their complement) with the Compliance Receive bit is set to 1b and the Loopback bit is set to 0b. Pseudo Port B is defined as the other Pseudo Port.
- The Retimer determines the highest common data rate of the Link by examining the data rate identifiers in the TS1 Ordered Sets received on each Pseudo Port, and the max data rate supported by the Retimer.
- If the highest common data rate is equal to 5.0 GT/s then:
  - The Retimer must change its data rate to 5.0 GT/s as described in Section 4.4.6.4.
  - The Retimer Pseudo Port A must set its de-emphasis according to the selectable de-emphasis bit received in the eight consecutive TS1 Ordered Sets.
  - The Retimer Pseudo Port B must set its de-emphasis in an implementation specific manner.
- If the highest common data rate is equal to 8.0 GT/s then:
  - The Retimer must change its data rate to 8.0 GT/s as described in Section 4.4.6.4.
  - Lane numbers are determined as described in Section 4.2.11.
  - The Retimer Pseudo Port A must set its transmitter coefficients on each Lane to the transmitter preset value advertised in Symbol 6 of the eight consecutive TS1 Ordered Sets and this value must be used by the transmitter (use of the receiver preset value advertised in those TS1 Ordered Sets is optional). If the common data rate is 8.0 GT/s, any Lanes that did not receive eight consecutive TS1 Ordered Sets with Transmitter Preset information can use any supported Transmitter Preset in an implementation specific manner.
  - The Retimer Pseudo Port B must set its Transmitter and Receiver equalization in an implementation specific manner.
- The Retimer must forward the Modified Compliance Pattern when it has locked to the pattern. This occurs independently on each Lane in each direction. If a Lane's Receiver

loses Symbol Lock or Block Alignment, the associated Transmitter (i.e., same Lane on opposite Pseudo Port) Continues to forward data.

- ❑ Once locked to the pattern, the Retimer keeps an internal count of received Symbol errors, on a per Lane basis. The pattern lock and Lane error is permitted to be readable in an implementation specific manner, on a per Lane basis.

- ❑ When operating with 128b/130b encoding, Symbols with errors are forwarded unmodified by default, or may optionally be corrected to remove error pollution. The default behavior must be supported and the method of selecting the optional behavior, if supported, is implementation specific.

- ❑ When operating with 8b/10b encoding, Symbols with errors are replaced with the D21.3 Symbol with incorrect disparity by default, or may optionally be corrected to remove error pollution. The default behavior must be supported and the method of selecting the optional behavior, if supported, is implementation specific.

- ❑ The error status Symbol when using 8b/10b encoding or the Error_Status field when using 128b/130b encoding is forwarded unmodified by default, or may optionally be redefined as it is transmitted by the Retimer. The default behavior must be supported and the method of selecting the optional behavior, if supported, is implementation specific.

- ❑ If any Lane receives an EIOS on either Pseudo Port then:

  - Transmit EIOS on every Lane of the Pseudo Port that did not receive EIOS and place the Transmitters in Electrical Idle. Place the Transmitters of the other Pseudo Port in Electrical Idle; EIOS is not transmitted by the other Pseudo Port.

  - Set the RT_next_data_rate variable to 2.5 GT/s.

  - Set the RT_error_data_rate variable to 2.5 GT/s.

  - The Compliance Receive additional rules no longer apply unless the rules for entering this Section are met again.

4.4.6.15 Enter Compliance Rules

The Retimer follows these additional rules if the Retimer is exiting Electrical Idle after entering Electrical Idle as a result of Hot Reset, and the Retimer Enter Compliance bit is set in the Retimer. The purpose of the following rules is to support Link operation with a Retimer when the Link partners enter compliance as a result of the Enter Compliance bit in the Link Control 2 Register set to 1b in both Link Components and a Hot Reset occurring on the Link. Retimers do not support Link operation if the Link partners enter compliance when they exit detect if the entry into detect was not caused by a Hot Reset.

Retimers must support the following register fields in an implementation specific manner:

- ❑ Retimer Target Link Speed

  - One field per Retimer

  - Type = RWS

- Size = 3 bits
- Default = 001b
- Encoding:
    - 001b = 2.5 GT/s
    - 010b = 5.0 GT/s
    - 011b = 8.0 GT/s

❑ Retimer Transmit Margin

- One field per Pseudo Port
- Type = RWS
- Size = 3 bits
- Default = 000b
- Encoding:
    - 000b = Normal Operating Range
    - 001b-111b = As defined in Section 4.3.3.6, not all encodings are required to be implemented

❑ Retimer Enter Compliance

- One bit per Retimer
- Type = RWS
- Size = 1 bit
- Default = 0b
- Encoding:
    - 0b = do not enter compliance
    - 1b = enter compliance

❑ Retimer Enter Modified Compliance

- One bit per Retimer
- Type = RWS
- Size = 1 bit
- Default = 0b

- Encoding:
    - 0b = do not enter modified compliance
    - 1b = enter modified compliance

❑ Retimer Compliance SOS

- One bit per Retimer
- Type = RWS
- Size = 1 bit
- Default = 0b
- Encoding:
    - 0b = Send no SKP Ordered Sets between sequences when sending the Compliance Pattern or Modified Compliance Pattern with 8b/10b encoding.
    - 1b = Send two SKP Ordered Sets between sequences when sending the Compliance Pattern or Modified Compliance Pattern with 8b/10b encoding.

❑ Retimer Compliance Preset/De-emphasis

- One field per Pseudo Port
- Type = RWS
- Size = 4 bits
- Default = 0000b
- Encoding when Retimer Target Link Speed is 5.0 GT/s:
    - 0000b -6.0 dB
    - 0001b -3.5 dB
- Encoding when Retimer Target Link Speed is 8.0 GT/s: the Transmitter Preset.

A Retimer must examine the values in the above registers when the Retimer exits from Hot Reset. If the Retimer Enter Compliance bit is set to 1b the following rules apply:

❑ The Retimer adjusts its data rate as defined by Retimer Target Link Speed. No data is forwarded until the data rate change has occurred.

❑ The Retimer configures its Transmitters according to Retimer Compliance Preset/De-emphasis on a per Pseudo Port basis.

❑ The Retimer must forward the Compliance or Modified Compliance Pattern when it has locked to the pattern. The Retimer must search for the Compliance Pattern if the Retimer

Enter Modified Compliance bit is set to 0b or search for the Modified Compliance Pattern if the Retimer Enter Modified Compliance bit is set to 1b. This occurs independently on each Lane in each direction.

- ❑ When using 8b/10b encoding, a particular Lane's Receiver independently determines a successful lock to the incoming Modified Compliance Pattern or Compliance Pattern by looking for any one occurrence of the Modified Compliance Pattern or Compliance Pattern.
  - An occurrence is defined above as the sequence of 8b/10b Symbols defined in Section 4.2.8.
  - In the case of the Modified Compliance Pattern, the error status Symbols are not to be used for the lock process since they are undefined at any given moment.
  - Lock must be achieved within 1.0 ms of receiving the Modified Compliance Pattern.
- ❑ When using 128b/130b encoding each Lane determines Pattern Lock independently when it achieves Block Alignment as described in Section 4.2.2.2.1.
  - Lock must be achieved within 1.5 ms of receiving the Modified Compliance Pattern or Compliance Pattern.
- ❑ When 128b/130b encoding is used, Symbols with errors are forwarded unmodified by default, or may optionally be corrected to remove error pollution. The default behavior must be supported and the method of selecting the optional behavior, if supported, is implementation specific.
- ❑ When 8b/10b encoding is used, Symbols with errors are replaced with the D21.3 Symbol with incorrect disparity by default, or may optionally be corrected to remove error pollution. The default behavior must be supported.
- ❑ Once locked, the Retimer keeps an internal count of received Symbol errors, on a per Lane basis. If the Retimer is forwarding the Modified Compliance Pattern then the error status Symbol when using 8b/10b encoding or the Error_Status field when using 128b/130b encoding is forwarded unmodified by default, or may optionally be redefined as it is transmitted by the Retimer. The default behavior must be supported and the method of selecting the optional behavior, if supported, is implementation specific. The Retimer is permitted to make the pattern lock and Lane error information available in an implementation specific manner, on a per Lane basis.
- ❑ If the data rate is not 2.5 GT/s and an EIOS is received on any Lane then:
  - All Lanes in that direction transmit 8 EIOS and then all Transmitters in that direction are placed in Electrical Idle.
  - When both directions have sent 8 EIOS and placed their Transmitters in Electrical Idle the data rate is changed to 2.5 GT/s.
  - Set the RT_next_data_rate variable to 2.5 GT/s.
  - Set the RT_error_data_rate variable to 2.5 GT/s.

- The Retimer Enter Compliance bit and Retimer Enter Modified Compliance bit are both set to 0b.
- The above additional rules no longer apply unless the rules for entering this Section and clause are met again.

❑ If the data rate is 2.5 GT/s and TS1 Ordered Sets are received then:

- The TS1 Ordered Sets are forwarded.
- The Retimer Enter Compliance bit and Retimer Enter Modified Compliance bit are both set to 0b.
- The above additional rules no longer apply unless the rules for entering this Section and clause are met again.

4.4.7 Execution Mode Rules

In Execution mode, Retimers directly control all information transmitted by the Pseudo Ports rather than forwarding information.

4.4.7.1 CompLoadBoard Rules

While the Retimer is in the CompLoadBoard (Compliance Load Board) state both Pseudo Ports are executing the protocol as regular Ports, generating Symbols as specified in the following sub-sections on each Port, rather than forwarding from one Pseudo Port to the other.

IMPLEMENTATION NOTE

Passive Load on Transmitter

This state is entered when a passive load is placed on one Pseudo Port, and the other Pseudo Port is receiving traffic.

4.4.7.1.1 CompLoadBoard.Entry

❑ RT_LinkUp = 0b.

❑ The Pseudo Port that received Compliance Pattern (Pseudo Port A) does the following:

- The data rate remains at 2.5 GT/s.
- The Transmitter is placed in Electrical Idle.
- The Receiver ignores incoming Symbols.

❑ The other Pseudo Port (Pseudo Port B) does the following:

- The data rate remains at 2.5 GT/s.

- The Transmitter is placed in Electrical Idle. Receiver detection is performed on all Lanes as described in Section 4.4.4.
- The Receiver ignores incoming Symbols.

❑ If Pseudo Port B's Receiver detection determines there are no Receivers attached on any Lanes, then the next state for both Pseudo Ports is CompLoadBoard.Exit.

❑ Else the next state for both Pseudo Ports is CompLoadBoard.Pattern.

4.4.7.1.2 CompLoadBoard.Pattern

When The Retimer enters CompLoadBoard.Pattern the following occur:

❑ Pseudo Port A does the following:

- The Transmitter remains in Electrical Idle.
- The Receiver ignores incoming Symbols.

❑ Pseudo Port B does the following:

- The Transmitter sends out the Compliance Pattern on all Lanes that detected a Receiver at the data rate and de-emphasis/preset level determined as described in Section 4.2.6.2.2, (i.e., each consecutive entry into Compliance.Pattern advances the pattern). If the new data rate is not 2.5 GT/s, the Transmitter is placed in Electrical Idle prior to the data rate change. The period of Electrical Idle must be greater than 1 ms but it is not to exceed 2 ms.

❑ If Pseudo Port B detects an Electrical Idle exit of any Lane that detected a Receiver, then the next state for both Pseudo Ports is CompLoadBoard.Exit.

4.4.7.1.3 CompLoadBoard.Exit

When The Retimer enters CompLoadBoard.Exit the following occur:

❑ The Pseudo Port A:

- Data rate remains at 2.5 GT/s.
- The Transmitter sends 32 K28.7 Symbols (EIE) on the Lane(s) where Electrical Idle exit was detected on Pseudo Port B. After the 32 Symbols the Transmitter is placed in Electrical Idle.
- The Receiver ignores incoming Symbols.

❑ Pseudo Port B:

- If the Transmitter is transmitting at a rate other than 2.5 GT/s the Transmitter sends eight consecutive EIOS.

- The Transmitter is placed in Electrical Idle. If the Transmitter was transmitting at a rate other than 2.5 GT/s the period of Electrical Idle must be at least 1.0 ms.
- Data rate is changed to 2.5 GT/s, if not already at 2.5 GT/s.

❑ Both Pseudo Ports are placed in Forwarding mode.

IMPLEMENTATION NOTE

TS1 Ordered Sets in Forwarding mode

Once in Forwarding mode one of two things will likely occur:

❑ TS1 Ordered Sets are received and forwarded from Pseudo Port's B Receiver to Pseudo Port's A Transmitter. Link training continues.

❑ Or: TS1 Ordered Sets are not received because 100 MHz pulses are being received on a lane from the compliance load board, advancing the Compliance Pattern. In this case the Retimer must transition from Forwarding mode to CompLoadBoard when the device attached to Pseudo Port A times out from Polling.Active to Polling.Compliance. The Retimer advances the Compliance Pattern on each entry to CompLoadBoard .

4.4.7.2 Link Equalization Rules

When in the Execution mode performing Link Equalization, the Pseudo Ports act as regular Ports, generating Symbols on each Port rather than forwarding from one Pseudo Port to the other. When the Retimer is in Execution mode it must use the Lane and Link numbers stored in RT_captured_lane_number and RT_captured_link_number.

This mode is entered while the Upstream and Downstream Components on the Link are in negotiation to enter Phase 2 of the Equalization procedure following the procedure for switching to Execution mode described in Section 4.4.5.

4.4.7.2.1 Downstream Lanes

The LF and FS values received in two consecutive TS1 Ordered Sets when the Upstream Port is in Phase 0 must be stored for use during Phase 3, if the Downstream Pseudo Port wants to adjust the Upstream Port's Transmitter.

4.4.7.2.1.1 Phase 2

Transmitter behaves as described in Section 4.2.6.4.2.1.2 except as follows:

❑ Next phase is Phase 3 Active if all configured Lanes receive two consecutive TS1 Ordered Sets with EC=11b.

- ❑ Else, next state is Force Timeout after a 32 ms timeout with a tolerance of -0 ms and +4 ms.

4.4.7.2.1.2 Phase 3 Active

Transmitter behaves as described in Section 4.2.6.4.2.1.3 except the 24 ms timeout is 2.5 ms and as follows:

- ❑ Next phase is Phase 3 Passive if all configured Lanes are operating at their optimal settings.
- ❑ Else, next state is Force Timeout after a timeout of 2.5 ms with a tolerance of -0 ms and +0.1 ms

4.4.7.2.1.3 Phase 3 Passive

- ❑ Transmitter sends TS1 Ordered Sets with EC = 11b, and the Transmitter Preset field and the Coefficients fields must not be changed from the final value transmitted in Phase 3 Active.
- ❑ The transmitter switches to Forwarding mode when the Upstream Pseudo Port exits Phase 3.

4.4.7.2.2 Upstream Lanes

The LF and FS values received in two consecutive TS1 Ordered Sets when the Downstream Port is in Phase 1 must be stored for use during Phase 2, if the Upstream Pseudo Port wants to adjust the Downstream Port's Transmitter.

4.4.7.2.2.1 Phase 2 Active

- ❑ Transmitter behaves as described in Section 4.2.6.4.2.2.3 except the 24 ms timeout is 2.5 ms and as follows:
  - Next state is Phase 2 Passive if all configured Lanes are operating at their optimal settings.
  - Else, next state is Force Timeout after a 2.5 ms timeout with a tolerance of -0 ms and +0.1 ms

4.4.7.2.2.2 Phase 2 Passive

- ❑ Transmitter sends TS1 Ordered Sets with EC = 10b, and the Transmitter Preset field and the Coefficients fields must not be changed from the final value transmitted in Phase 2 Active.
- ❑ The next state is Phase 3 when the Downstream Pseudo Port has completed Phase 3 active.

4.4.7.2.2.3 Phase 3

Transmitter follows Phase 3 rules for Upstream Lanes in Section 4.2.6.4.2.2.4 except as follows:

- If all configured Lanes receive two consecutive TS1 Ordered Sets with EC=00b then the Retimer switches to Forwarding mode.
- Else, next state is Force Timeout after a timeout of 32 ms with a tolerance of -0 ms and +4 ms

4.4.7.2.3 Force Timeout

- The Modified Compliance Pattern is transmitted by both Pseudo Ports at 8.0 GT/s for a minimum of 1.0 ms.
- If a Receiver receives an EIOS or infers Electrical Idle (inferred Electrical Idle is via not detecting an exit from Electrical Idle see Table 4-b) on any Lane the Transmitters on all lanes of the opposite Pseudo Port stop sending the Modified Compliance Pattern and are placed in Electrical Idle after sending an EIOS.
- When both Paths have placed their Transmitters in Electrical Idle, the Retimer changes its data rate back to the previous data rate, on both Pseudo Ports, and enters Forwarding mode.
- The Transmitters of both Pseudo Ports must be in Electrical Idle for at least 6 us, before forwarding data.

IMPLEMENTATION NOTE

Purpose of Force Timeout State

The purpose of this state is to ensure both Link Components are in Recovery.Speed at the same time so they go back to the previous data rate.

4.4.7.3 Slave Loopback

Retimers optionally support Slave Loopback in Execution mode. By default Retimers are configured to forward loopback between loopback master and loopback slave. Retimers are permitted to allow configuration in an implementation specific manner to act as a loopback slave on either Pseudo Port. The other Pseudo Port that is not the loopback slave, places its Transmitter in Electrical Idle, and ignores any data on its Receivers.

4.4.7.3.1 Slave Loopback.Entry

The Pseudo Port that did not receive the TS1 Ordered Set with the Loopback bit set to 1b does the following:

- The Transmitter is placed in Electrical Idle.

- The Receiver ignores incoming Symbols.

The Pseudo Port that did receive the TS1 Ordered Set with the Loopback bit set to 1b behaves as the loopback slave as described in Section 4.2.6.10.1 with the following exceptions:

- The statement "LinkUp = 0b (False)" is replaced by "RT_LinkUp = 0b".
- The statement "If Loopback.Entry was entered from Configuration.Linkwidth.Start" is replaced by "If SlaveLoopback.Entry was entered when RT_LinkUp =0b".
- References to Loopback.Active become Slave Loopback.Active.

4.4.7.3.2 Slave Loopback.Active

The Pseudo Port that did not receive the TS1 Ordered Set with the Loopback bit set to 1b does the following:

- The Transmitter remains in Electrical Idle.
- The Receiver continues to ignore incoming Symbols.

The Pseudo Port that did receive the TS1 Ordered Set with the Loopback bit set to 1b behaves as the loopback slave as described in Section 4.2.6.10.2 with the following exception:

- References to Loopback.Exit become Slave Loopback.Exit.

4.4.7.3.3 Slave Loopback.Exit

The Pseudo Port that did not receive the TS1 Ordered Set with the Loopback bit set to 1b must do the following:

- Maintain the Transmitter in Electrical Idle.
- Set the data rate to 2.5 GT/s.
- The Receiver continues to ignore incoming Symbols.

The Pseudo Port that did receive the TS1 or TS2 Ordered Set with the Loopback bit set to 1b must behave as the loopback slave as described in Section 4.2.6.10.3 with the following exception:

- The clause "The next state of the loopback master and loopback slave is Detect" becomes "The Data rate is set to 2.5 GT/s and then both Pseudo Ports are placed in Forwarding mode.".

4.4.8 Retimer Latency

This Section defines the requirements on allowed Retimer Latency.

4.4.8.1 Measurement

Latency must be measured when the Retimer is in Forwarding mode and the Link is in L0, and is defined as the time from when the last bit of a Symbol is received at the input pins of one Pseudo Port to when the equivalent bit is transmitted on the output pins of the other Pseudo Port.

Retimer vendors are strongly encouraged to specify the latency of the Retimer in their data sheets.

Retimers are permitted to have different latencies at different data rates, and when this is the case it is strongly recommended the latency be specified per data rate.

4.4.8.2 Maximum Limit on Retimer Latency

Retimer latency shall be less than the following limit, when not operating in SRIS.

Table 1-c: Retimer Latency Limit not SRIS (Symbol times)

	2.5 GT/s	5.0 GT/s	8.0 GT/s
Maximum Latency	32	32	64

4.4.8.3 Impacts on Upstream and Downstream Components

Retimers will add to the channel latency. The round trip delay is 4 times the specified latency when two Retimers are present. It is recommended that designers of Upstream and Downstream Components consider Retimer latency when determining the following characteristics:

- Data Link Layer Retry Buffer size
- Transaction Layer Receiver buffer size and Flow Control Credits
- Data Link Layer REPLAY_TIMER Limits

Additional buffering (replay or FC) may be required to compensate for the additional channel latency.

4.4.9 SRIS

Retimers are permitted but not required to support SRIS. Retimers that support SRIS must provide a mechanism for enabling the higher rate of SKP Ordered Set transmission, as Retimers must generate SKP Ordered Sets while in Execution mode. Retimers that are enabled to support SRIS will incur additional latency in the elastic store between receive and transmit clock domains. The additional latency is required to handle the case where a Max_Payload_Size TLP is transmitted and SKP Ordered Sets, which are scheduled, are not sent. The additional latency is a function of Link width and Max_Payload_Size. This additional latency is not included in Table 4-c.

A SRIS capable Retimer must provide an implementation specific mechanism to configure the supported Max_Payload_Size while in SRIS, that must be configured to be greater than or equal to the Max_Payload_Size for the Transmitter in the Port that the Pseudo Port is receiving. Retimer latency must be less than the following limit for the current supported Max_Payload_Size, with SRIS.

Table 4-d: Retimer Latency Limit SRIS (Symbol times)

MaxPayload Size	2.5 GT/s	5.0 GT/s	8.0 GT/s
128 Bytes	34 (max)	34 (max)	66 (max)
256 Bytes	36 (max)	36 (max)	68 (max)
512 Bytes	39 (max)	39 (max)	71 (max)
1024 Bytes	46 (max)	46 (max)	78 (max)
2048 Bytes	59 (max)	59 (max)	91 (max)
4096 Bytes	86 (max)	86 (max)	118 (max)

4.4.9.1 L1 PM Substates Support

The following Section describes the Retimer's requirements to support the optional L1 PM Substates.

The Retimer enters L1.1 when CLKREQ# is sampled as de-asserted. The following occur:

- ❑ REFCLK to the Retimer is turned off.
- ❑ The PHY remains powered.
- ❑ The Retimer places all Transmitters in Electrical Idle on both Pseudo Ports (if not already in Electrical Idle, the expected state). Transmitters maintain their common mode voltage.
- ❑ The Retimer must ignore any Electrical Idle exit from all Receivers on both Pseudo Ports.

The Retimer exits L1.1 when CLKREQ# is sampled as asserted. The following occur:

- ❑ REFCLK to the Retimer is enabled.
- ❑ Normal operation of the Electrical Idle exit circuit is resumed on all Lanes of both Pseudo Ports of the Retimer.
- ❑ Normal exit from Electrical Idle exit behavior is resumed, See 4.4.6.3.

Retimers do not support L1.2, but if they support L1.1 and the removal of the reference clock then they must not interfere with the attached components ability to enter L1.2.

Retimer vendors must document specific implementation requirements applying to CLKREQ#. For example, a Retimer implementation that does not support the removal of the reference clock might require an implementation to pull CLKREQ# low.

IMPLEMENTATION NOTE

CLKREQ# Connection Topology with a Retimer Supporting L1 PM Substates

In this platform configuration the Upstream Component (A) has only a single CLKREQ# signal. The Upstream and Downstream Ports' CLKREQ# (A and C), and the Retimer's CLKREQB# signals are connected to each other. In this case, Upstream Component (A), must assert CLKREQ# signal whenever it requires a reference clock. Component A, Component B, and the Retimer have their REFCLKs removed/restored at the same time.

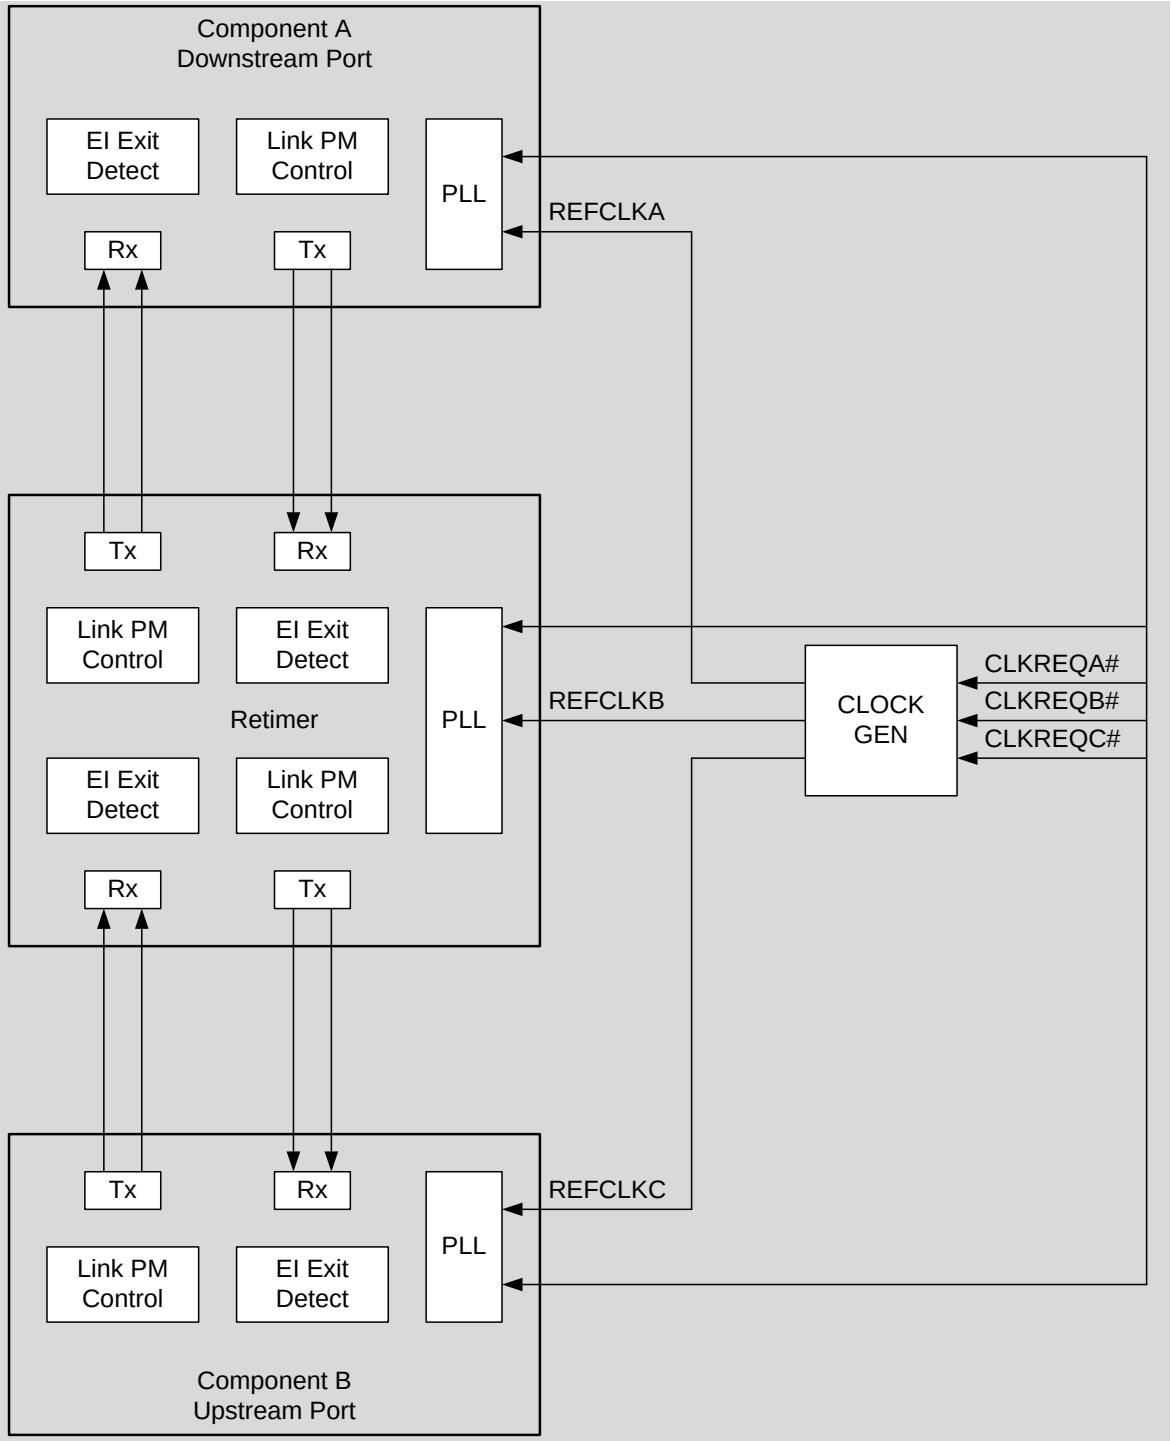

Figure 4-b Retimer CLKREQ# Connection Topology

4.4.10 Retimer Configuration Parameters

Retimers must provide an implementation specific mechanism to configure each of the parameters in this Section.

The parameters are split into two groups: parameters that are configurable globally for the Retimer and parameters that are configurable for each physical Retimer Pseudo Port.

If a per Pseudo Port parameter only applies to an Upstream or a Downstream Pseudo Port the Retimer is not required to provide an implementation specific mechanism to configure the parameter for the other type of Pseudo Port.

4.4.10.1 Global Parameters

- **Port Orientation Method**. This controls whether the Port Orientation is determined dynamically as described in Section 4.4.6.2, or statically based on vendor assignment of Upstream and Downstream Pseudo Ports. If the Port Orientation is set to static the Retimer is not required to dynamically adjust the Port Orientation as described in Section 4.4.6.2. The default behavior is for the Port Orientation to be dynamically determined.

- **Maximum Data Rate.** This controls the maximum data rate that the Retimer sets in the Data Rate Identifier field of training sets that the Retimer transmits. Retimers that support only the 2.5 GT/s speed are permitted not to provide this configuration parameter.

- **SRIS Enable.** This controls whether the Retimer is configured for SRIS and transmits SKP ordered sets at the SRIS mode rate when in Execution mode. Retimers that do not support SRIS and at least one other clocking architecture are not required to provide this configuration parameter.

- **SRIS Link Payload Size.** This controls the maximum payload size the Retimer supports while in SRIS. The value must be selectable from all the Maximum Payload Sizes shown in Table 4-d. The default value of this parameter is to support a payload size of 4096 bytes. Retimers that do not support SRIS are not required to provide this configuration parameter.

  The following are example of cases where it might be appropriate to configure the SRIS Link Payload Size to a smaller value than the default:

  - A Retimer is part of a motherboard with a Root Port that supports a maximum payload size less than 4096 bytes.
  - A Retimer is part of an add-in card with an Endpoint that supports a Maximum Payload Size less than 4096 bytes.
  - A Retimer is located Downstream of the Downstream Port of a Switch integrated as part of a system, the Root Port silicon supports a Maximum Payload Size less than 4096 bytes and the system does not support peer to peer traffic.

4.4.10.2 Per Physical Pseudo Port Parameters

- **Port Orientation**. This is applicable only when the Port Orientation Method is configured for static determination. This is set for either Upstream or Downstream.

Each Pseudo Port must be configured for a different orientation, or the behavior is undefined.

- **Selectable De-emphasis**. When the Downstream Pseudo Port is operating at 5.0 GT/s this controls the transmit de-emphasis of the Link to either -3.5 dB or -6 dB in specific situations and the value of the Selectable De-emphasis field in training sets transmitted by the Downstream Pseudo Port. See Sections 4.2.6 for detailed usage information. When the Link Segment is not operating at the 5.0 GT/s speed, the setting of this bit has no effect. Retimers that support only the 2.5 GT/s speed are permitted not to provide this configuration parameter.

- **Rx Impedance Control.** This controls whether the Retimer dynamically applies and removes 50 Ω terminations or statically has 50 Ω terminations present. The value must be selectable from Dynamic, Off, and On. The default behavior is Dynamic.

- **Tx Compliance Disable.** This controls whether the Retimer transmits the Compliance Pattern in the CompLoadBoard.Pattern state. The default behavior is for the Retimer to transmit the Compliance Pattern in the CompLoadBoard.Pattern state. If TX Compliance Pattern is set to disabled, the Retimer Transmitters remain in Electrical Idle and do not transmit Compliance Pattern in CompLoadBoard.Pattern – all other behavior in the CompLoadBoard state is the same.

- **Pseudo Port Slave Loopback.** This controls whether the Retimer operates in a Forwarding mode during loopback on the Link or enters Slave Loopback on the Pseudo Port. The default behavior is for the Retimer to operate in Forwarding mode during loopback. Retimers that do not support optional Slave Loopback are permitted not to provide this configuration parameter. This configuration parameter shall only be enabled for one physical Port. Retimer behavior is undefined if the parameter is enabled for more than one physical Port.

- **Downstream Pseudo Port TX Preset.** This controls the initial TX preset used by the Downstream Pseudo Port transmitter for 8.0 GT/s transmission. The default value is implementation specific. The value must be selectable from all applicable values in Table 4-3.

- **Downstream Pseudo Port Requested TX Preset.** This controls the initial transmitter preset value used in the training sets transmitted by the Downstream Pseudo Port. The default value is implementation specific. The value must be selectable from all values in Table 4-3.

- **Downstream Pseudo Port RX Hint.** This controls the Receiver Preset Hint value used in the training sets transmitted by the Downstream Pseudo Port. The default value is implementation specific. The value must be selectable from all values in Table 4-4.

Modify Figure 7-27

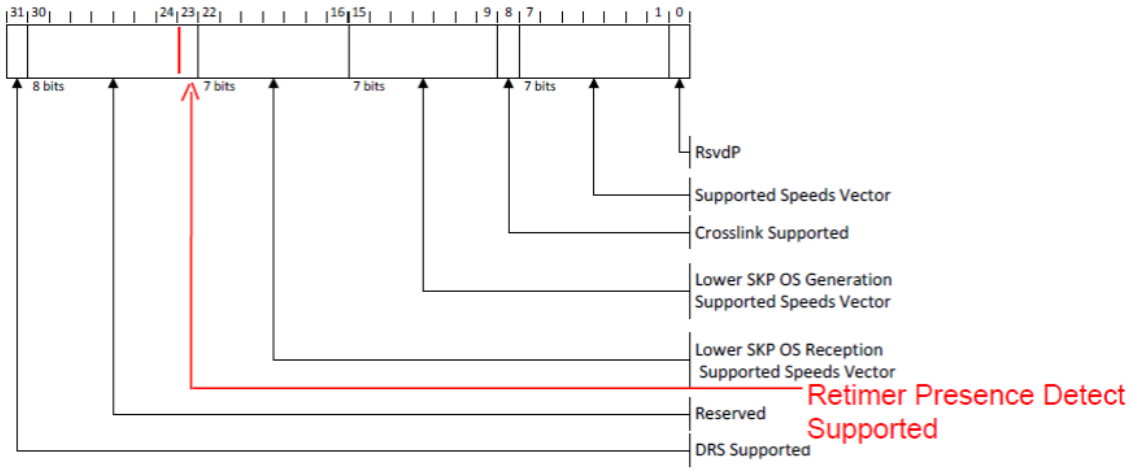

Figure 7-27: Link Capabilities 2 Register

Modify Section 7.8.18 Link Capabilities 2 Register, Table

…		
23	**Retimer Presence Detect Supported** – For a Downstream Port, when set to 1b, this bit indicates that the associated Port supports detection and reporting of Retimer presence. This bit is valid for Downstream Ports. This bit is Reserved for Upstream Ports.	HwInit/RsvdP

Modify Figure 7-29

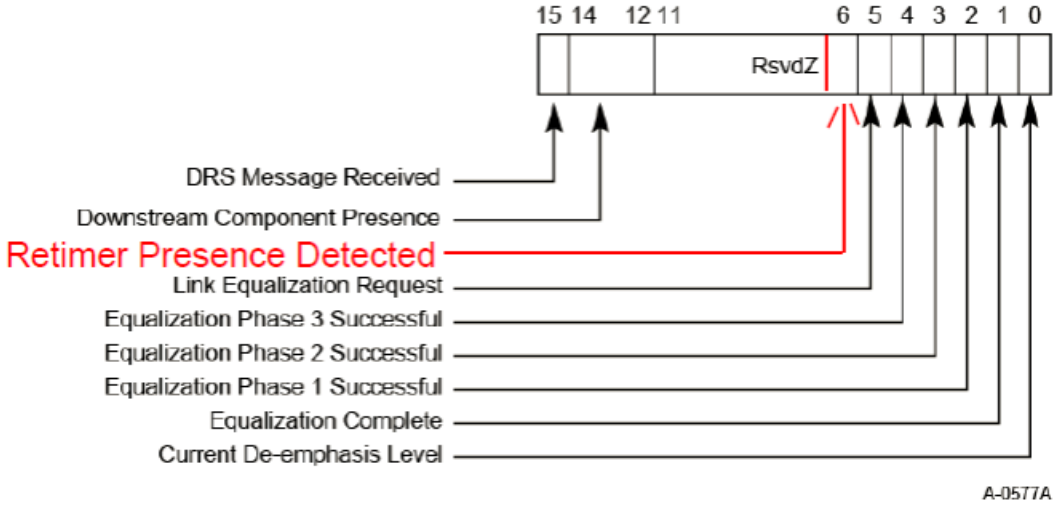

Figure 7-29: Link Status 2 Register

Modify Section 7.8.20 Link Status 2 Register,Table

…		
6	**Retimer** Presence Detected – For a Downstream Port, when set to 1b, this bit indicates that a Retimer was present during the most recent Link negotiation. Refer to Section 4.2.6.3.5.1 for details. The default value of this bit is 0b. This bit is required for Downstream Ports that have the Retimer Presence Detect Supported bit Set. Downstream Ports that have the Retimer Presence Detect Supported bit of the Link Capabilities 2 register set to 0b are permitted to hardwire this bit to 0b. This bit is Reserved for Upstream Ports.	ROS/RsvdZ

Modify Table 4-66 as shown:

Link Capabilities 2	Retimer Presence Detected Supported	If Port Mapped to Bridge is Set, returns value associated with the Port. If Port Mapped to Bridge is Clear, returns 1b if any Switch Port supports the feature.	Returns the value associated with the Port
Link Control 2	Target Link Speed Enter Compliance Hardware Autonomous Speed Disable Selectable De-emphasis Transmit Margin Enter Modified Compliance Compliance SOS Compliance De-emphasis	Read/Write registers, value is ignored.	Base
Link Status 2	Current De-emphasis Level Retimer Presence Detected	Contains the value associated with the Port.	Base