Compute Express Link™ (CXL™)

Engineering Change Notice to the Specification 2.0

September 2021

QoS Telemetry Compliance Testcases

LEGAL NOTICE FOR THIS PUBLICLY-AVAILABLE SPECIFICATION FROM COMPUTE EXPRESS LINK CONSORTIUM, INC.

This CXL **Specification** (this "**CXL Specification**" or this "**document**") is owned by and is proprietary to Compute Express Link Consortium, Inc., a Delaware nonprofit corporation (sometimes referred to as "**CXL**" or the "**CXL Consortium**" or the "**Company**") and/or its successors and assigns.

NOTICE TO USERS WHO ARE MEMBERS OF THE CXL CONSORTIUM:

If you are a Member of the CXL Consortium (sometimes referred to as a "**CXL Member**"), and even if you have received this publicly-available version of this CXL Specification after agreeing to CXL Consortium's Evaluation Copy Agreement (a copy of which is available https://www.computeexpresslink.org/download-the-specification, each such CXL Member must also be in compliance with all of the following CXL Consortium documents, policies and/or procedures (collectively, the "**CXL Governing Documents**") in order for such CXL Member's use and/or implementation of this CXL Specification to receive and enjoy all of the rights, benefits, privileges and protections of CXL Consortium membership: (i) CXL Consortium's Intellectual Property Policy; (ii) CXL Consortium's Bylaws; (iii) any and all other CXL Consortium policies and procedures; and (iv) the CXL Member's Participation Agreement.

NOTICE TO NON-MEMBERS OF THE CXL CONSORTIUM:

If you are **not** a CXL Member and have received this publicly-available version of this CXL Specification, your use of this document is subject to your compliance with, and is limited by, all of the terms and conditions of the CXL Consortium's Evaluation Copy Agreement (a copy of which is available at https://www.computeexpresslink.org/download-the-specification**).**

In addition to the restrictions set forth in the CXL Consortium's Evaluation Copy Agreement, any references or citations to this document must acknowledge the Compute Express Link Consortium, Inc.'s sole and exclusive copyright ownership of this CXL Specification. The proper copyright citation or reference is as follows: **"© 2019-2021 COMPUTE EXPRESS LINK CONSORTIUM, INC. ALL RIGHTS RESERVED."** When making any such citation or reference to this document you are not permitted to revise, alter, modify, make any derivatives of, or otherwise amend the referenced portion of this document in any way without the prior express written permission of the Compute Express Link Consortium, Inc.

Except for the limited rights explicitly given to a non-CXL Member pursuant to the explicit provisions of the CXL Consortium's Evaluation Copy Agreement which governs the publicly-available version of this CXL Specification, nothing contained in this CXL Specification shall be deemed as granting (either expressly or impliedly) to any party that is not a CXL Member: (ii) any kind of license to implement or use this CXL Specification or any portion or content described or contained therein, or any kind of license in or to any other intellectual property owned or controlled by the CXL Consortium, including without limitation any trademarks of the CXL Consortium.; or (ii) any benefits and/or rights as a CXL Member under any CXL Governing Documents.

LEGAL DISCLAIMERS FOR ALL PARTIES:

THIS DOCUMENT AND ALL SPECIFICATIONS AND/OR OTHER CONTENT PROVIDED HEREIN IS PROVIDED ON AN "**AS IS**" BASIS. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, COMPUTE EXPRESS LINK CONSORTIUM, INC. (ALONG WITH THE CONTRIBUTORS TO THIS DOCUMENT) HEREBY DISCLAIM ALL REPRESENTATIONS, WARRANTIES AND/OR COVENANTS, EITHER EXPRESS OR IMPLIED, STATUTORY OR AT COMMON LAW, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, VALIDITY, AND/OR NON-INFRINGEMENT.
In the event this CXL Specification makes any references (including without limitation any incorporation by reference) to another standard's setting organization's or any other party's ("**Third Party**") content or work, including without limitation any specifications or standards of any such Third Party ("**Third Party Specification**"), you are hereby notified that your use or implementation of any Third Party Specification: (i) is not governed by any of the CXL Governing Documents; (ii) may require your use of a Third Party's patents, copyrights or other intellectual property rights, which in turn may require you to independently obtain a license or other consent from that Third Party in order to have full rights to implement or use that Third Party Specification; and/or (iii) may be governed by the intellectual property policy or other policies or procedures of the Third Party which owns the Third Party Specification. Any trademarks or service marks of any Third Party which may be referenced in this CXL Specification is owned by the respective owner of such marks.

NOTICE TO ALL PARTIES REGARDING THE PCI-SIG UNIQUE VALUE PROVIDED IN THIS CXL SPECIFICATION:

NOTICE TO USERS: THE UNIQUE VALUE THAT IS PROVIDED IN THIS CXL SPECIFICATION IS FOR USE IN VENDOR DEFINED MESSAGE FIELDS, DESIGNATED VENDOR SPECIFIC EXTENDED CAPABILITIES, AND ALTERNATE PROTOCOL NEGOTIATION ONLY AND MAY NOT BE USED IN ANY OTHER MANNER, AND A USER OF THE UNIQUE VALUE MAY NOT USE THE UNIQUE VALUE IN A MANNER THAT (A) ALTERS, MODIFIES, HARMS OR DAMAGES THE TECHNICAL FUNCTIONING, SAFETY OR SECURITY OF THE PCI-SIG ECOSYSTEM OR ANY PORTION THEREOF, OR (B) COULD OR WOULD REASONABLY BE DETERMINED TO ALTER, MODIFY, HARM OR DAMAGE THE TECHNICAL FUNCTIONING, SAFETY OR SECURITY OF THE PCI-SIG ECOSYSTEM OR ANY PORTION THEREOF (FOR PURPOSES OF THIS NOTICE, "**PCI-SIG ECOSYSTEM**" MEANS THE PCI-SIG SPECIFICATIONS, MEMBERS OF PCI-SIG AND THEIR ASSOCIATED PRODUCTS AND SERVICES THAT INCORPORATE ALL OR A PORTION OF A PCI-SIG SPECIFICATION AND EXTENDS TO THOSE PRODUCTS AND SERVICES INTERFACING WITH PCI-SIG MEMBER PRODUCTS AND SERVICES).

CXL ENGINEERING CHANGE NOTICE

TITLE:	QoS Telemetry Compliance Testcases
DATE:	Introduced date (03/05/2021) Updated date (05/24/2021)
AFFECTED DOCUMENT:	CXL 2.0 Specification
SPONSOR:	Luis E. Rodriguez (Avery) Compliance Workgroup

Part I

1. **Summary of Functional Changes**

   QoS telemetry was not included as part of the initial CXL 2.0 Compliance Chapter. The testcases proposed in this ECN allows for basic testing of QoS Telemetry controls as well as the FM API and SLD Control and Status commands. The proposed changes are only an add-on to the existing compliance program and do not require any hardware changes in either the host test system or devices.

2. **Benefits as a Result of the Changes**
   Allows for the Compliance Testing program to ensure interoperability and basic functionality of QoS Telemetry as well as the FM API and SLD Control and Status command sets.

3. **Assessment of the Impact**
   Compliance program software would have to implement the new testcases.

4. **Analysis of the Hardware Implications**
   No hardware changes required.

5. **Analysis of the Software Implications**
   No system software changes required for host or device systems.

6. **Analysis of the Compliance and Test Implications**
   Compliance program would have to implement the new testcases as part of the CV tool or any other future compliance testing solutions.

Part II

Detailed Description of the change

In the Compliance Chapter, section 14.3 add the following new subsection defining two new testcases for QoS Compliance testing:

14.3.6.1.5 Egress Port Backpressure Test

This test applies to an MLD that supports FM API or an SLD that supports the Memory Device command set. This test sets up the Device to execute Algorithm 1a, 1b and 2 in succession in order to stress data path for CXL.mem transactions. An equivalent version of the algorithm is setup to be executed by **Host** software so as to enable false sharing of the cachelines. Test system software and configuration details are determined by the host platform and are system specific. NUMBER_OF_QOS_TEST_LOOPS, NUMBER_OF_CHECK_AVERAGE, and BackpressureSample Interval setting in the test steps below is decided upon by the testing platform/software.

Test Steps:

For an MLD:

. Through the FM API, check if Egress Port Congestion Supported is set by issuing a Get LD Info command.

. If Egress Port Congestion Supported is enabled:

Repeat for NUMBER_OF_QOS_TEST_LOOPS:

. Set the BackpressureSample Interval setting to a value between 1 -31 through the Set QoS Control command

. Set the Egress Port Congestion Enable bit through the Set QoS Control command

. Check that the Egress Port Congestion Enable bit was set successfully in the Get QoS Control Response.

. Run equivalent of Algorithm 1a, 1b and 2 in succession on the Host and Device targeting device attached memory.
. While Algorithms 1a, 1b and 2 are running: Check the reported Backpressure Average Percentage through the Get QoS Status command and response. It should report values within the valid range which is 0 – 100. Repeat this step NUMBER_OF_CHECK_AVERAGE times at a certain interval.

For an SLD:

. Check if Egress Port Congestion Supported is set by issuing an Identify Memory Device, and checking the corresponding Identify Memory Device Output Payload

. If Egress Port Congestion Supported is enabled:

Repeat for NUMBER_OF_QOS_TEST_LOOPS:

. Set the BackpressureSample Interval setting to a value between 1 -31 through the Set SLD QoS Control Request command

. Set the Egress Port Congestion Enable bit through the Set SLD QoS Control Request

. Check that the Egress Port Congestion Enable bit was set successfully in the Get SLD QoS Control Response.

. Check the reported Backpressure Average Percentage through the Get QoS Status command and response

. Run the equivalent of Algorithms 1a, 1b and 2 in succession on the Host and Device targeting device attached memory.

. While Algorithms 1a, 1b and 2 are running: Check the reported Backpressure Average Percentage through the Get SLD QoS Status command and response. It should report values within the valid range which is 0 – 100. Repeat this step NUMBER_OF_CHECK_AVERAGE times at a certain interval.

Pass Criteria:

. Egress Port Congestion Enable is set after enabling it

. Backpressure Average Percentage reports valid values within 0 -100.

. No data corruptions or system errors reported while executing Algorithm 1a, 1b, 2

Fail Criteria:

. Egress Port Congestion Enable is set not after enabling it

. Backpressure Average Percentage reports any value outside of the valid 0 -100 range.

. Data corruptions or system errors reported while executing Algorithm 1a, 1b, 2

14.3.6.1.6 Temporary Throughput Reduction Test

This test applies to an MLD that supports FM API or an SLD that supports the Memory Device Command set. This test sets up the Device to execute Algorithm 1a, 1b and 2 in succession in order to stress data path for CXL.mem transactions. An equivalent version of the algorithm is setup to be executed by **Host** software so as to enable false sharing of the cachelines. Test system software and configuration details are determined by the host platform and are system specific. NUMBER_OF_QOS_TEST_LOOPS in the test steps is decided upon by the testing platform/software.

Test Steps:

For an MLD:

. Through the FM API, check if Temporary Throughput Reduction Supported is set by issuing a Get LD Info command.

. If Temporary Throughput Reduction Supported is enabled:

Repeat for NUMBER_OF_QOS_TEST_LOOPS:

. Set the Temporary Throughput Reduction Enable bit by issuing the Set QoS Control command

. Check that the Temporary Throughput Reduction Enable bit was set successfully in the Get QoS Control Response.

. Run the equivalent of Algorithm 1a, 1b and 2 in succession on the Host and Device targeting device attached memory.

For an SLD:

. Through the Memory Device Command set, check if Temporary Throughput Reduction Supported is set by issuing an Identify Memory Device, and checking corresponding Identify Memory Device Output Payload

. If Temporary Throughput Reduction Supported is enabled:

Repeat for NUMBER_OF_QOS_TEST_LOOPS:

. Set the Temporary Throughput Reduction Enable bit through the Set SLD QoS Control Request

. Check that the Temporary Throughput Reduction Enable bit was set successfully in the Get SLD QoS Control Response.

. Run the equivalent of Algorithm 1a, 1b and 2 in succession on the Host and Device targeting device attached memory.

Pass Criteria:

. Temporary Throughput Reduction Enable is set after enabling it

. No data corruptions or system errors reported while executing Algorithm 1a, 1b, 2

Fail Criteria:

. Temporary Throughput Reduction Enable is set not after enabling it

. Data corruptions or system errors reported while executing Algorithm 1a, 1b, 2