

UNH-IOL NVMe Testing Service

Test Plan for NVMe PCIe Transport Conformance
Version 23.0
Target Specification: NVMe over PCIe Transport
Specification 1.0c
Technical Document



Last Updated: January 23, 2025

UNH-IOL NVMe Testing Service
21 Madbury Rd Suite 100
Durham, NH 03824

Tel: +1 603-862-0090
Fax: +1 603-862-4181
Email: nvmelab@iol.unh.edu

TABLE OF CONTENTS

TABLE OF CONTENTS	2
MODIFICATION RECORD	3
ACKNOWLEDGMENTS	4
INTRODUCTION	5
REFERENCES	7
ABBREVIATIONS	8
Group 1: Controller Architecture	9
Test 1.1 – Controller Level Reset – Conventional Reset (M)	10
Test 1.2 – Controller Level Reset – Function Level Reset (M)	11
Test 1.3 – Controller Level Reset – NVM Subsystem Reset (M)	12
Test 1.4 – Controller Level Reset – Offset Values Persistent Across Reset (FYI)	14
Group 2: Power State Transitions	15
Test 2.1 – Autonomous Power State Transitions Enabled (M)	16
Test 2.2 – Return from Non-Operational State (FYI)	17
Case 1: Basic Operation (IP)	17
Case 2: Non-Operation State Admin Commands (IP)	17
Test 2.3 – Autonomous Power State Transition (M)	19
Case 1: Proper Structure	19
Test 2.4 – Power State Entrance Latency (M)	20
Test 2.5 – Power State Exit Latency (M)	21
Test 2.6 – Relative Read Throughput (M)	22
Test 2.7 – Relative Write Throughput (M)	23
Test 2.8 – Host Controlled Thermal Management (M)	24
Case 1: Basic Operation (M)	24
Case 2: Invalid Field (M)	24
Group 3: System Bus Registers (PCIeC Group 10)	26
Test 3.1 – PCI Express Capability Registers (M)	27
Group 4: Physical Interface Receiver Eye Opening Measurement	28
Test 4.1 – Physical Interface Eye Opening Measurement Log Page	29
Case 1: Physical Interface Receiver Eye Opening Measurement Log Page 19h Expected Values(FYI)	29
Case 2: Physical Interface Receiver Eye Opening Measurement Log Page 19h Size Based on EOMIP Value (FYI)	29
Case 3: Physical Interface Receiver Eye Opening Measurement Log Page 19h Lane Descriptor Order (FYI)	30
Case 4: Physical Interface Receiver Eye Opening Measurement Log Page 19h EOM Header (FYI)	30
Appendix A: DEFAULT TEST SETUP	32
Appendix B: NOTES ON TEST PROCEDURES	33
Appendix C: TEST TOOLS	34
Appendix D: NVME INTEGRATORS LIST REQUIREMENTS	35

MODIFICATION RECORD

2022 January 21 (Version 1.0) Initial Release

Tim Sheehan:

1. Initial release of the NVMe PCIe Transport Conformance Test Suite v17, inclusive of PCIe tests from the NVMe Conformance Test Suite v16.0

2022 July 14 (Version 18.0) Final Release

Tim Sheehan:

1. Program Revision Update

2023 January 04 (Version 19.0) Final Release

Tim Sheehan:

Tests with Status Changes

1. Test 2.2.1 updated from (FYI) to (M)
2. Test 2.2.2 updated from (FYI) to (M)

2023 July 19 (Version 20.0) Final Release

Tim Sheehan:

Tests with Status Changes

1. Test 2.2.1 & 2 updated to (IP), incorrectly changed during last test plan modification.
2. Modified test case 2.6.1 and 2.7.1 procedures and results steps for Power State Transition

July 26, 2024 (Version 22.0)

Carter Snay:

New and Modified test cases:

1. Modified test case 4.1.4 by removing the requirement for the Printable Eye bit to be set to 1, and added a step to read the Physical Interface Receiver Eye Opening Measurement log page between the start measurement and aborting measurement steps.

2025 January 23 (Version 23.0)

Carter Snay:

New and Modified test cases

1. Updated test 2.7.1 to check if APST is enabled, and if it is, disable it prior to the operations in the test to avoid interference.

ACKNOWLEDGMENTS

The UNH-IOL would like to acknowledge the efforts of the following individuals in the development of this test plan:

Tim Sheehan
Carter Snay

UNH InterOperability Laboratory
UNH InterOperability Laboratory

INTRODUCTION

The University of New Hampshire’s InterOperability Laboratory (IOL) is an institution designed to improve the interoperability of standards-based products by providing a neutral environment where a product can be tested against other implementations of a common standard, both in terms of interoperability and conformance. This particular suite of tests has been developed to help implementers evaluate the NVMe functionality of their products. This test suite is aimed at validating products in support of the work being directed by the NVMe Promoters Group.

These tests are designed to determine if a product conforms to specifications defined in the NVM Express PCIe Transport Specification Revision 1.0 specification and NVM Express Base Specification 2.0a. Successful completion of these tests provide a reasonable level of confidence that the Device Under Test (DUT) will function properly in many NVMe environments. Transport testing should also include those tests for tests in respective I/O command set test plans

The tests contained in this document are organized in order to simplify the identification of information related to a test, and to facilitate in the actual testing process. Tests are separated into groups, primarily in order to reduce setup time in the lab environment, however the different groups typically also tend to focus on specific aspects of device functionality. A two-number, dot-notated naming system is used to catalog the tests. This format allows for the addition of future tests in the appropriate groups without requiring the renumbering of the subsequent tests. The NVMe 2.0 refactoring effort has made it important to reference pre-2.0 test cases that have moved to a 2.0 refactored test plan. This new test case numbers will be realized at the end of each test case name with the legacy test case number noted.

The test definitions themselves are intended to provide a high-level description of the motivation, resources, procedures, and methodologies specific to each test. Formally, each test description contains the following sections:

Purpose

The purpose is a brief statement outlining what the test attempts to achieve. The test is written at the functional level.

References

This section specifies all reference material *external* to the test suite, including the specific references for the test in question, and any other references that might be helpful in understanding the test methodology and/or test results. External sources are always referenced by a bracketed number (e.g., [1]) when mentioned in the test description. Any other references in the test description that are not indicated in this manner refer to elements within the test suite document itself (e.g., “Appendix 5.A”, or “Table 5.1.1-1”).

Resource Requirements

The requirements section specifies the test hardware and/or software needed to perform the test. This is generally expressed in terms of minimum requirements, however in some cases specific equipment manufacturer/model information may be provided.

Last Modification

This specifies the date of the last modification to this test.

Discussion

The discussion covers the assumptions made in the design or implementation of the test, as well as known limitations. Other items specific to the test are covered here as well.

Test Setup

The setup section describes the initial configuration of the test environment. Small changes in the configuration should not be included here, and are generally covered in the test procedure section (next).

Procedure

The procedure section of the test description contains the systematic instructions for carrying out the test. It provides a cookbook approach to testing, and may be interspersed with observable results. These procedures should be the ideal test methodology, independent of specific tool limitations or restrictions.

Observable Results

This section lists the specific observable items that can be examined by the tester in order to verify that the DUT is operating properly. When multiple values for an observable are possible, this section provides a short discussion on how to interpret them. The determination of a pass or fail outcome for a particular test is generally based on the successful (or unsuccessful) detection of a specific observable.

Possible Problems

This section contains a description of known issues with the test procedure, which may affect test results in certain situations. It may also refer the reader to test suite appendices and/or other external sources that may provide more detail regarding these issues.

REFERENCES

The following documents are referenced in this text:

Old Ref : NVM Express Revision 1.4 Specification (June 10, 2019)

1. NVM Express Base Specification Revision 2.0c (October 4, 2022)
2. NVMe over PCIe Transport Specification version 1.0c (October 3, 2022)

ABBREVIATIONS

The following abbreviations are applied to the test titles of each of the tests described in this document for indicating the status of test requirements.

M - Mandatory

FYI - FYI

IP - In Progress

The following abbreviations applied to the test titles of each of the tests described in this document for indicating what product types a test may apply to. It is assumed that all tests apply to base NVMe products using PCIe.

OF – Test applies to NVMe-oF products

The following legacy numbering is applied to the test titles of each of the test cases to map the origin of the test case movement due to 2.0 refactoring efforts. If there is no change to the test case number then this field will be left off of the 2.0 refactored test case title.

NC – UNH-IOL NVMe Conformance Test Plan version 16.0, September 23, 2021

Legacy numbering : (PCIeC:x.x)

Group 1: Controller Architecture

Overview:

This section describes a method for performing conformance verification for NVMe products implementing the NVMe Controller Architecture.

Notes:

The preliminary draft descriptions for the tests defined in this group are considered complete, and the tests are pending implementation (during which time additional revisions/modifications are likely to occur).

Test 1.1 – Controller Level Reset – Conventional Reset (M)

Purpose: To verify that an NVMe Controller performs the proper actions when a Conventional Reset occurs.

References:

- Old Ref : NVMe Specification 7.3.2
- NVM Express Base specification 2.0a : 3.7.2
- NVMe over PCIe Transport Specification 1.0a : 3.3.1

Resource Requirements:

Tools capable of monitoring and decoding traffic on the NVMe interface.

Last Modification: November 26, 2018

Discussion: When a Controller Level Reset occurs, the Host and Controller are required to take specific action.

Test Setup: See Appendix A.

Test Procedure:

1. Configure the NVMe Device to perform a Conventional Controller Level Reset. Repeat this for each of the following cases:
 - a. PCI Express Hot reset
2. When the reset is complete, configure the NVMe Host to issue a Write and then a Read command to the NVMe Controller.

Observable Results:

1. Verify that the NVMe Controller is able to properly execute NVMe Write and Read commands after the reset is complete.
2. Verify that the NVMe Controller performs the following actions when each reset case defined above is initiated:
 - a. The controller stops processing any outstanding Admin or I/O commands.
 - b. All I/O Submission Queues are deleted.
 - c. All I/O Completion Queues are deleted.
 - d. The controller is brought to an Idle state. When this is complete, CSTS.RDY is cleared to '0'.
 - e. All controller registers defined in section 3 of the NVMe Specification and internal controller state are reset.

Possible Problems: None.

Test 1.2 – Controller Level Reset – Function Level Reset (M)

Purpose: To verify that an NVMe Controller performs the proper actions when a Function Level Reset occurs.

References:

Old Ref : NVMe Specification 7.3.2
NVMe over PCIe Transport Specification 1.0a : 3.3.1

Resource Requirements:

Tools capable of monitoring and decoding traffic on the NVMe interface.

Last Modification: March 16, 2016

Discussion: When a Controller Level Reset occurs, the Host and Controller are required to take specific action.

Support for the Function Level Reset mechanism is indicated in the Function Level Reset Capability (FLRC) field (bit 28) of the PCI Express Device Capabilities (PXDCAP) (offset PXCAP + 4h) PCI Express Register. All NVMe Controllers must support Function Level Reset and so this register value should be set to '1' for all NVMe Controllers.

A Function Level Reset is initiated by the NVMe Host by writing a value of '1' to the Initiate Function Level Reset field (bit 15) of the PCI Express Device Control (PXDC) (offset PXCAP + 8h) PCI Express Register. The value read by software from this field shall always be '0'.

Test Setup: See Appendix A.

Test Procedure:

1. Configure the NVMe Host to read the PXDCAP.FLRC PCI Express register field.
2. Configure the NVMe Host to read the Initiate Function Level Reset field of the PXDC PCI Express register.
3. Configure the NVMe Host to write a value of '1' to the Initiate Function Level Reset field of the PXDC PCI Express register in order to initiate a Function Level Reset.
4. When the reset is complete, configure the NVMe Host to issue a Write command to the NVMe Controller.

Observable Results:

1. Verify that the value read from the PXDCAP.FLRC PCI Express register field is '1' to indicate support for the Function Level Reset mechanism.
2. Verify that the value read from the Initiate Function Level Reset field of the PXDC PCI Express register is '0'.
3. Verify that the NVMe Controller is able to properly execute the Write command after the reset is complete.
4. Verify that the NVMe Controller performs the following actions when the Function Level Reset is initiated:
 - a. The controller stops processing any outstanding Admin or I/O commands.
 - b. All I/O Submission Queues are deleted.
 - c. All I/O Completion Queues are deleted.
 - d. The controller is brought to an Idle state. When this is complete, CSTS.RDY is cleared to '0'.
 - e. All controller registers defined in section 3 of the NVMe Specification and internal controller state are reset.

Possible Problems: A reliable means of performing this test has not been determined. Therefore, this test should not be included in any industry approved determination of conformance.

Test 1.3 – Controller Level Reset – NVM Subsystem Reset (M)

Purpose: To verify that an NVMe Controller performs the proper actions when an NVM Subsystem Reset occurs.

References:

Old Ref : NVMe Specification 7.3.1 and 7.3.2
NVMe over PCIe Transport Specification 1.0a : 3.3.1

Resource Requirements:

Tools capable of monitoring and decoding traffic on the NVMe interface.

Last Modification: April 2, 2019

Discussion: When a Controller Level Reset occurs, the Host and Controller are required to take specific action.

An NVM Subsystem Reset is initiated when:

1. Power is applied to the NVM System,
2. A value of 4E564D65h (“NVMe”) is written to the NSSR.NSSRC controller register field, or
3. A vendor specific event occurs.

When an NVM Subsystem Reset occurs, the entire NVM subsystem is reset. This includes the initiation of a Controller Level Reset on all controllers that make up the NVM subsystem and a transition to the Detect LTSSM state by all PCI Express ports of the NVM subsystem.

The occurrence of an NVM Subsystem Reset while power is applied to the NVM subsystem is reported by the initial value of the CSTS.NSSRO field following the NVM Subsystem Reset. This field may be used by host software to determine if the sudden loss of communication with a controller was due to an NVM Subsystem Reset or some other condition.

This test is only applicable if the CAP.NSSRS field is set ‘1’ to indicate support for writing to the NSSR.NSSRC field.

Test Setup: See Appendix A.

Test Procedure:

1. Configure the NVMe Host to read the CAP.NSSRS field to determine if the DUT supports writing to the NSSR.NSSRC field. If the NSSRS field is cleared to ‘0’, the test is not applicable. If the CAP.NSSRS field is set to ‘1’, continue to step 2.
2. Configure the NVMe host to read the CSTS.NSSRO register value, and record the value. If the value is not zero, perform a power cycle of the DUT and restart the test.
3. Configure the NVMe host to write a value of 4E564D65h (“NVMe”) to the NSSR.NSSRC field.
4. When the reset is complete, the PCIe link is reestablished, the NVMe controller is enabled and an Identify is performed.

Observable Results:

1. Verify that when the reset is performed, the drive is temporarily no longer visible from the host system. Depending on the reset and bring up speed of the device, this may not be observable.
2. Verify that the NVMe Controller is able to properly execute NVMe Identify command after the reset is complete.
3. Verify that the NVMe Controller performs the following actions when the NVM Subsystem Reset is initiated:
 - a. The controller stops processing any outstanding Admin or I/O commands.
 - b. All I/O Submission Queues are deleted.
 - c. All I/O Completion Queues are deleted.
 - d. The controller is brought to an Idle state. When this is complete, CSTS.RDY is cleared to ‘0’.

- e. All controller registers defined in section 3 of the NVMe Specification and internal controller state are reset.

Possible Problems: The DUT may or may not support NVM Subsystem Reset as it is an optional NVMe feature. When the reset is performed, the drive is temporarily no longer visible from the host system. Depending on the reset and bring up speed of the device, this may not be observable.

Test 1.4 – Controller Level Reset – Offset Values Persistent Across Reset (FYI)

Purpose : In TP4100 PMR Controller Reset Enhancements, there are new requirements set for what is no longer allowed to be reset during a Controller Level Reset. This test ensures that the values are not reset as is now expected by the 2.0 NVMe Base Specification.

References:

Old Ref : NVMe Specification 7.3.1 and 7.3.2
NVMe over PCIe Transport Specification 1.0a : 3.3.1

Resource Requirements:

Tools capable of monitoring and decoding traffic on the NVMe interface.

Last Modification: January 18, 2024

Test Setup: See Appendix A.

Test Procedure:

1. Configure the NVMe Host to read the following offsets and save the values for later comparison.
 - a. Offset 24h: AQA – Admin Queue Attributes
 - b. Offset 28h: ASQ – Admin Submission Queue Base Address
 - c. Offset 30h: ACQ – Admin Completion Queue Base Address
 - d. Offset E00h: PMRCAP – Persistent Memory Region Capabilities
 - e. Offset E04h: PMRCTL – Persistent Memory Region Control
 - f. Offset E08h: PMRSTS – Persistent Memory Region Status
 - g. Offset E0Ch: PMREBS – Persistent Memory Region Elasticity Buffer Size
 - h. Offset E10h: PMRSWTP – Persistent Memory Region Sustained Write Throughput
 - i. Offset 50h: CMBMSC – Controller Memory Buffer Memory Space Control
2. Configure the NVMe Host to send a Controller Level Reset
3. Configure the NVMe host to re-read the register values listed in step 1.

Observable Results:

1. Verify that each offset value is identical to the value before the Controller Level Reset.

Possible Problems: The DUT may or may not support NVM Subsystem Reset as it is an optional NVMe feature. When the reset is performed, the drive is temporarily no longer visible from the host system. Depending on the reset and bring up speed of the device, this may not be observable.

Group 2: Power State Transitions

Overview:

This section describes a method for performing conformance verification for NVMe products implementing NVMe Power State Transitions.

Notes:

The preliminary draft descriptions for the tests defined in this group are considered complete, and the tests are pending implementation (during which time additional revisions/modifications are likely to occur).

Test 2.1 – Autonomous Power State Transitions Enabled (M)

Purpose: To determine if an NVMe Controller properly supports Autonomous Power State Transitions.

References:

Old Ref : NVMe Specification 5.14.1.12, Fig 90, Fig 108, Fig 124
NVM Express Base specification 2.0a : 5.27.1.19
NVMe over PCIe Transport Specification 1.0a : 3.6 & 3.8.2.3

Resource Requirements:

Tools capable of monitoring and decoding traffic on the NVMe interface.

Last Modification: April 13, 2015

Discussion: An NVMe device may change power states autonomously without host intervention if Autonomous Power State Transitions are supported by the device and enabled by the Host. There are 32 allowable power states defined sequentially in the 256 byte data structure entry.

Each entry is 64 bits long and describes the Idle Time Prior to Transition (ITPTT) and Idle Transition Power State (ITPS). The Idle Transition Power State specifies the next power state the device will transition to after there is a continuous period of idle time in the current power state that exceeds the time specified in the Idle Time Prior to Transition field.

This test is not applicable to devices that do not claim to support Autonomous Power State Transitions.

Test Setup: See Appendix A.

Test Procedure:

1. Check that the DUT supports Autonomous Power State Transitions by setting Bit 0 of Byte 265 of the Identify Controller Data Structure, to 1. If this bit is set to 0, the test is not performed.
2. Using the Set Features Command, enable Feature Identifier 0Ch for Autonomous Power State Transition.
3. Perform the Get Feature Command to see that the Autonomous Power State Transitions feature was enabled (APSTE), and receive the Autonomous Power State Transition data structure.

Observable Results:

1. Verify that the controller returns a properly formatted Autonomous Power State Transition data structure.

Possible Problems: A reliable means of performing this test has not been determined. Therefore, this test should not be included in any industry approved determination of conformance.

Test 2.2 – Return from Non-Operational State (FYI)

Purpose: To determine if an NVMe Controller properly supports Autonomous Power State Transitions.

References:

- Old Ref : NVMe Specification 8.4.1
- NVM Express Base specification 2.0a : 5.27.1.19
- NVMe over PCIe Transport Specification 1.0a : 3.6 & 3.8.2.3

Resource Requirements:

Tools capable of monitoring and decoding traffic on the NVMe interface.

Last Modification: January 23, 2018

Discussion: When in a non-operational power state, regardless of whether autonomous power state transitions are enabled, the controller shall autonomously transition back to the last operational power state when an I/O Submission Queue Tail Doorbell is written.

When in a non-operational power state, regardless of whether autonomous power state transitions are enabled, the controller shall not transition back if an Admin Command has been sent.

Test Setup: See Appendix A.

Test Setup: See Appendix A.

Case 1: Basic Operation (IP)

Test Procedure:

1. Check that the DUT supports Autonomous Power State Transitions by setting Bit 0 of Byte 265 of the Identify Controller Data Structure, to 1. If this bit is set to 0, the test is not performed.
2. Using the Set Features Command, enable Feature Identifier 0Ch for Autonomous Power State Transition.
3. Perform the Get Feature Command to see that the Autonomous Power State Transitions feature was enabled (APSTE), and receive the Autonomous Power State Transition data structure.
4. If the Autonomous Power State Transition Data Structure indicates that the device supports entering a non-operational state via APST, allow the DUT to remain idle for ITPT for each power state successively until the DUT enters a non-operational state.
5. Perform Identify Power State Descriptor Data Structure, check the NOPS field.
6. Perform an NVMe I/O Command, such as NVMe Write to the DUT.
7. Perform Identify Power State Descriptor Data Structure, check the NOPS field.

Observable Results:

1. Verify that the controller returns a properly formatted Autonomous Power State Transition data structure.
2. Verify that after the writing of the I/O Submission Queue Tail Doorbell, through the write command, the DUT returns to the last operational power state.

Case 2: Non-Operation State Admin Commands (IP)

Test Procedure:

1. Check that the DUT supports Autonomous Power State Transitions by setting Bit 0 of Byte 265 of the Identify Controller Data Structure, to 1. If this bit is set to 0, the test is not performed.
2. Using the Set Features Command, enable Feature Identifier 0Ch for Autonomous Power State Transition.
3. Perform the Get Feature Command to see that the Autonomous Power State Transitions feature was enabled (APSTE), and receive the Autonomous Power State Transition data structure.

4. If the Autonomous Power State Transition Data Structure indicates that the device supports entering a non-operational state via APST, allow the DUT to remain idle for ITPT for each power state successively until the DUT enters a non-operational state.
5. For each supported Admin Command opcode, preform an Admin Command of that opcode, then perform Identify Power State Descriptor Data Structure, check the NOPS field.

Observable Results:

1. Verify that the controller returns a properly formatted Autonomous Power State Transition data structure.
2. Verify that after each Admin Command, the DUT stays in the Non-Operational Power State.

Possible Problems: None.

Test 2.3 – Autonomous Power State Transition (M)

Purpose: To verify that an NVMe system can properly handle autonomous power state transitions.

References:

Old Ref : NVMe Specification 8.4.2, 5.14.1.12
NVM Express Base specification 2.0a : 5.27.1.19
NVMe over PCIe Transport Specification 1.0a : 3.6 & 3.8.2.3

Resource Requirements:

Tools capable of monitoring and decoding traffic on the NVMe interface.

Last Modification: April 1, 2016

Discussion: The controller may support autonomous power state transitions, as indicated in the Identify Controller data structure at byte 265. Autonomous power state transitions provide a mechanism for the host to configure the controller to automatically transition between power states on certain conditions without software intervention.

The entry condition to transition to the Idle Transition Power State is that the controller has been in idle for a continuous period of time exceeding the Idle Time Prior to Transition time specified. The controller is idle when there are no commands outstanding to any I/O Submission Queue. The power state to transition to shall be a non-operational power state (a non-operational power state may autonomously transition to another non-operational power state). If an operational power state is specified then the controller should abort the command with a status of Invalid Field in Command.

Test Setup: See Appendix A

Case 1: Proper Structure

Discussion: Each entry in the Autonomous Power State Transition data structure is defined in the NVMe Specification. Each entry is 64 bits in size. There is an entry for each of the allowable 32 power states. For power states that are not supported, the unused Autonomous Power State Transition data structure entries shall be cleared to all zeroes. The entries begin with power state 0 and then increase sequentially (i.e., power state 0 is described in bytes 7:0, power state 1 is described in bytes 15:8, etc.). The data structure is 256 bytes in size and shall be physically contiguous.

Test Procedure:

1. Configure the NVMe Host to issue a Set Features command with an unsupported power state.
2. Configure the NVMe Host to issue a Get Features command with the feature ID set to Autonomous Power State Transition (0Ch) to the NVMe Controller.

Observable Results:

1. Verify that after the completion of the command, the controller posts a completion queue entry to the Admin Completion Queue indicating the status for the command.
2. Verify that the Autonomous Power State Transition data structure and each of its entries are of proper size.
3. Verify that Autonomous Power State Transition data structure entries for power states not supported by the controller are cleared to all zeroes.

Test 2.4 – Power State Entrance Latency (M)

Purpose: To verify that an NVMe system properly documents its entrance latency for each Power State.

Reference:

Old Ref : NVMe Specification 8.4
NVM Express Base specification 2.0a : 8.15
NVMe over PCIe Transport Specification 1.0a : 3.6

Resource Requirements:

Tools capable of monitoring and decoding traffic on the NVMe interface.

Last Modification: January 23, 2018

Discussion

The controller may support a Power States Entrance Latency field, as indicated in the Power State Descriptor table described in the NVMe specification. This field indicates how long it should take, in microseconds, for a controller's power state to transition to another.

Test Setup: See Appendix A

Test Procedure:

1. For each Power State in the Power State Descriptor table with an ENLAT field that is not 0h:
2. Record the ENLAT field of the selected Power State and the EXLAT field of the current Power State
3. Have the NVMe host send a Set Feature Command with Feature Identifier 02h, Power Management, indicating the new Power State
4. Wait the EXLAT of the previous Power State and the ENLAT of the power state being selected.
5. Send a Get Features command with Feature Identifier 02, Power Management.

Observable Results:

1. Verify that the Get Features command returns the Power State selected in the Set Features command.

Possible Problems: None.

Test 2.5 – Power State Exit Latency (M)

Purpose: To verify that an NVMe system properly documents its exit latency for each Power State.

Reference:

Old Ref : NVMe Specification 8.4
NVM Express Base specification 2.0a : 8.15
NVMe over PCIe Transport Specification 1.0a : 3.6

Resource Requirements:

Tools capable of monitoring and decoding traffic on the NVMe interface.

Last Modification: January 23, 2018

Discussion

The controller may support a Power States Exit Latency field, as indicated in the Power State Descriptor table described in the NVMe specification. This field indicates how long it should take, in microseconds, for a controller's power state to transition to another.

Test Setup: See Appendix A

Test Procedure:

1. For each Power State in the Power State Descriptor table with an EXLAT field that is not 0h:
2. Record the ENLAT field of the selected Power State and the EXLAT field of the current Power State
3. Have the NVMe host send a Set Feature Command with Feature Identifier 02h, Power Management, indicating the new Power State
4. Wait the EXLAT of the previous Power State and the ENLAT of the power state being selected.
5. Send a Get Features command with Feature Identifier 02, Power Management.

Observable Results:

1. Verify that the Get Features command returns the Power State selected in the Set Features command.

Possible Problems: None.

Test 2.6 – Relative Read Throughput (M)

Purpose: To determine if an NVMe Controller properly supports the Relative Read Throughput rates supported by different power states correctly.

Reference:

Old Ref : NVMe Specification 8.4
NVM Express Base specification 2.0a : 8.15
NVMe over PCIe Transport Specification 1.0a : 3.6

Resource Requirements:

Tools capable of monitoring and decoding traffic on the NVMe interface.

Last Modification: January 23, 2018

Discussion

Different Power States have different Relative Read Throughput values that determine the speed at which they can read from the NVMe.

Test Setup: See Appendix A

Test Procedure:

1. Configure the host to perform 100 read operations for each power state and record the times.

Observable Results:

1. Verify the success of the 100 reads for each power state..

Possible Problems: None.

Test 2.7 – Relative Write Throughput (M)

Purpose: To determine if an NVMe Controller properly supports the Relative Write Throughput rates supported by different power states correctly.

Reference:

Old Ref : NVMe Specification 8.4
NVM Express Base specification 2.0a : 8.15
NVMe over PCIe Transport Specification 1.0a : 3.6

Resource Requirements:

Tools capable of monitoring and decoding traffic on the NVMe interface.

Last Modification: January 23, 2018

Discussion

Different Power States have different Relative Write Throughput values that determine the speed at which they can write to the NVMe.

Test Setup: See Appendix A

Test Procedure:

1. Configure the NVMe Host to issue a Get Features command with FID 0Ch and record the APSTE bit. If this bit is set to 1b, issue a Set Features command to set APSTE to 0, disabling APST.
2. Configure the host to perform 100 write operations for each power state and record the times.
3. If in step 1 APSTE was required to be cleared to 0, reenable it via a Set Features command.

Observable Results:

1. Verify the success of the 100 writes for each power state.
2. Verify the Set and Get Features commands complete successfully.

Possible Problems: None.

Test 2.8 – Host Controlled Thermal Management (M)

Purpose: To determine if an NVMe Controller properly supports the Host Controlled Thermal Management feature correctly.

Reference:

Old Ref : NVMe Specification 8.4
NVM Express Base Specification 2.0a : 5.27.1.13, 8.15.5

Resource Requirements:

Tools capable of monitoring and decoding traffic on the NVMe interface.

Last Modification: November 26, 2018

Discussion

The Host Controlled Thermal Management command can be set within certain boundaries specified by the Minimum Host Thermal Management Temperature field and the Maximum Host Thermal Management field in the Identify Controller Structure, this test ensures that the behavior of a Set Features command with an FID specifying Host Controlled Thermal Management reacts correctly to improper Thermal Management 1 & 2 values.

Test Setup: See Appendix A

Case 1: Basic Operation (M)

Test Procedure:

1. Record the current temperature of the DUT.
2. Perform an Identify Controller Data Structure (CNS=01h) to the DUT. Check the HCTMA (Host Controlled Thermal Management Attribute) bit. If HCTMA is set to 0, this test is not applicable. If HCTMA is set to 1, proceed to the next step.
3. Set the drive to the highest active power state
4. If possible, send a Set Features with FID 10h (Host Controlled Thermal Management), with a Thermal Management Temperature 1 that is below the current device temperature, and a Thermal Management Temperature 2 that is above the current temperature.
5. Record the current power state.
6. Set the drive to the highest active power state
7. If possible, send a Set Features with FID 10h (Host Controlled Thermal Management), with a Thermal Management Temperature 1 that is below the current device temperature, and a Thermal Management Temperature 2 that is above the Thermal Management Temperature 1, but below the current temperature.
8. Record the current power state.

Observable Results:

1. Verify that after each Set Features command completes correctly.
2. For informational purposes display the Power State after each Set Features command.

Case 2: Invalid Field (M)

Test Procedure:

1. Record the current power state.
2. Perform an Identify Controller Data Structure (CNS=01h) to the DUT. Check the HCTMA (Host Controlled Thermal Management Attribute) bit. If HCTMA is set to 0, this test is not applicable. If HCTMA is set to 1, proceed to the next step.
3. Perform a Set Features with FID 10h (Host Controlled Thermal Management), with a Thermal Management Temperature 1 that is 1 degree above the allowed, and a Thermal Management Temperature 2 that is 0xFFFF.

4. Record the current power state.
5. Perform a Set Features with FID 10h (Host Controlled Thermal Management), with a Thermal Management Temperature 1 that is 0xFFFF, and a Thermal Management Temperature 2 that is 1 degree above the Minimum Thermal Management Temperature.
6. Record the current power state.
7. Perform a Set Features with FID 10h (Host Controlled Thermal Management), with a Thermal Management Temperature 1 that is 1 degree above the allowed threshold, and a Thermal Management Temperature 2 that is 0x1.
8. Record the current power state.
9. Perform a Set Features with FID 10h (Host Controlled Thermal Management), with a Thermal Management Temperature 1 that is 0x1, and a Thermal Management Temperature 2 that is 1 degree above the Minimum Thermal Management Temperature.
10. Record the current power state.

Observable Results:

1. Verify that after each Set Features command, an error code of Invalid Field is indicated and the power state has not changed.

Possible Problems: None.

Group 3: System Bus Registers (PCIeC Group 10)

Overview:

This section describes a method for performing conformance verification for NVMe products implementing the System Bus Register.

Notes:

The preliminary draft descriptions for the tests defined in this group are considered complete, and the tests are pending implementation (during which time additional revisions/modifications are likely to occur).

Test 3.1 – PCI Express Capability Registers (M)

Purpose: To determine if an NVMe Controller properly implements the PCI Express Capability Registers.

References:

Old Ref : NVMe Specification 2.5
NVMe over PCIe Transport Specification 1.0a : 3.8.5

Resource Requirements:

Tools capable of monitoring and decoding traffic on the NVMe interface.

Last Modification: June 13, 2016

Discussion: NVMe controller using the PCI Express system bus must report their PCI Express capabilities via the PCI Express Capability Registers.

Test Setup: See Appendix A.

Test Procedure:

1. Configure the Host and Controller to bring up the PCIe link, and bring the NVMe Controller to the enabled state.
2. Examine the PCIe bringup sequence for the PXCAP register contents, and record. The PCIe capability structure ID is 0x10.

Observable Results:

1. Verify that the PCIe Capabilities register (PXCAP) follows the format defined in section 2.5 of the NVMe Specification.
2. Verify that all reserved bits are set to 0.

Possible Problems: None.

Group 4: Physical Interface Receiver Eye Opening Measurement

Overview:

This section describes a method for performing conformance verification for NVMe products implementing the Physical Interface Receiver Eye Opening Measurement Log Page

Notes:

The preliminary draft descriptions for the tests defined in this group are considered complete, and the tests are pending implementation (during which time additional revisions/modifications are likely to occur).

Test 4.1 – Physical Interface Eye Opening Measurement Log Page

Purpose: To determine if an NVMe Controller properly implements the PCI Express Capability Registers.

References:

NVMe over PCIe Transport Specification 1.0d

Resource Requirements:

Tools capable of monitoring and decoding traffic on the NVMe interface.

Last Modification: January 18, 2024

Case 1: Physical Interface Receiver Eye Opening Measurement Log Page 19h Expected Values (FYI)

Purpose: To verify proper implementation of newly defined bytes in the Physical Interface Receiver Eye Opening Measurement Log Page 19h

Test Procedure:

1. Configure the NVMe Host to issue a Get Log Page command for LID 00h
2. Read Bytes 79:76, if cleared to zero, this log page is not supported, and this test is not applicable.
3. Configure the NVMe Host to issue a Get Log Page command for LID 19h with the Log Specific Parameter Field (bits 11:10) set to 01b (Start Measurement and Read Log Data) and the Target Controller set to the current controller (bits 15:0) in the Log Specific Identifier Field.
4. Read the returned log page.

Observable Results:

1. Verify that the Get Log Page command completed successfully.
2. Verify that Byte 0 (Log Identifier) is set to 19h.
3. Verify that Byte 1 (EOM in Progress) is set to 1h.
4. Verify that Byte 2 (HSIZE) is set to 64.
5. Verify that Byte 9 (Log Revision) is set to 2h.
6. Verify that Byte 13 (Log Specific Parameter Field Copy) is set to 01b.
7. Verify that Byte 14 bits 3:0 are equal to the value of the Current Link Speed in the PXLS Register
8. Verify that Byte 18:19 is set to the Action field value of 10b

Case 2: Physical Interface Receiver Eye Opening Measurement Log Page 19h Size Based on EOMIP Value (FYI)

Purpose: To verify the Physical Interface Receiver Eye Opening Measurement Log Page is the expected size based on figure 72 of the NVM Express PCI Express Transport Specification version 1.0d

Test Procedure:

1. Configure the NVMe Host to issue a Get Log Page command for LID 00h
2. Read Bytes 79:76, if cleared to zero, this log page is not supported, and this test is not applicable.
3. Configure the NVMe Host to issue a Get Log Page command for LID 19h with the Log Specific Parameter Field (bits 11:10) set to 00b (Read Log Data) and the Target Controller set to the current controller (bits 15:0) in the Log Specific Identifier Field. Record the value of the Header Size (HSIZE) field (byte 3:2)
4. Configure the NVMe Host to issue a Get Log Page command for LID 19h with the Log Specific Parameter Field (bits 11:10) set to 01b (Start Measurement and Read Log Data) and the Target Controller set to the

current controller (bits 15:0) in the Log Specific Identifier Field. Record the value of the Header Size (HSIZE) field (byte 3:2)

5. Configure the NVMe Host to issue a Get Log Page command for LID 19h with the Log Specific Parameter Field (bits 11:10) set to 01b (Start Measurement and Read Log Data) and then immediately issue a second Get Log Page Command for LID 19 with the Log Specific Parameter field set to 10b (Abort Measurement and Clear Log) and the Target Controller set to the current controller (bits 15:0) in the Log Specific Identifier Field. Record the value of the Header Size (HSIZE) field (byte 3:2), Descriptor Size (DS) field (bytes 23:20), and Number of Descriptors (ND) field (bytes 25:24).

Observable Results:

1. Verify that all Get Log Page commands return successfully
2. Verify that the Log Page returned in step 3 and step 4 are 64 bytes in length.
3. Verify that the second Log Page returned in step 5 is has a length in bytes equal to Header Size plus the product of the Descriptor Size and Number of Descriptors fields (i.e. HSIZE + DS * ND)

Case 3: Physical Interface Receiver Eye Opening Measurement Log Page 19h Lane Descriptor Order (FYI)

Purpose: To verify the Physical Interface Receiver Eye Opening Measurement Log Page EOM Lane Descriptor are listed in the correct ascending order.

Test Procedure:

1. Configure the NVMe Host to issue a Get Log Page command for LID 00h
2. Read Bytes 79:76, if cleared to zero, this log page is not supported, and this test is not applicable.
3. Configure the NVMe Host to issue a Get Log Page command for LID 19h with the Log Specific Parameter Field (bits 11:10) set to 00b (Read Log Data) and the Target Controller set to the current controller (bits 15:0) in the Log Specific Identifier Field.

Observable Results:

1. Verify that all Get Log Page commands return successfully
2. Verify the EOM Lane Descriptors are listed in ascending order with the Lane number (byte 2) and Eye number (byte 3) of each Descriptor incrementally increases by 1

Case 4: Physical Interface Receiver Eye Opening Measurement Log Page 19h EOM Header (FYI)

Purpose: To verify the Physical Interface Receiver Eye Opening Measurement Log Page EOM Lane Descriptor are listed in the correct ascending order.

Test Procedure:

1. Configure the NVMe Host to issue a Get Log Page command for LID 00h
2. Read Bytes 79:76, if cleared to zero, this log page is not supported, and this test is not applicable.
3. Configure the NVMe Host to issue a Get Log Page command for LID 19h with the Log Specific Parameter Field (bits 11:10) set to 00b (Read Log Data) and the Target Controller set to the current controller (bits 15:0) in the Log Specific Identifier Field.
4. Configure the NVMe Host to issue a Get Log Page command for LID 19h with the Log Specific Parameter Field (bits 11:10) set to 01b (Start Measurement and Read Log Page Data) and the Target Controller set to the current controller (bits 15:0) in the Log Specific Identifier Field.
5. Issue the command from the previous step a second time
6. Configure the NVMe Host to issue a Get Log Page command for LID 19h with the Log Specific Parameter Field (bits 11:10) set to 00b (Read Log Data) and the Target Controller set to the current controller (bits 15:0) in the Log Specific Identifier Field.

7. Configure the NVMe Host to issue a Get Log Page command for LID 19h with the Log Specific Parameter Field (bits 11:10) set to 10b (Abort Measurement and Clear Log) and the Target Controller set to the current controller (bits 15:0) in the Log Specific Identifier Field.


Observable Results:

1. Verify that all Get Log Page commands return successfully
2. Verify the following values are present in the EOM Header:
 - a. Log ID is 19h
 - b. Header Size (HSIZE) is set to 64
 - c. EOM Data Generation Number is incremented by 1 between steps 4 and 5.
 - d. Log Revision field is set to 2h
 - e. Verify that if EOMIP = 2h, then Optional Data Present field is not set to 0h
 - f. Verify that Lanes and Eyes Per Lane are non-zero values
 - g. Number of Descriptors (ND) cleared to zero if EOMIP field is cleared to 0'
 - h. If Printable Eye (bit 0 of byte 10) is 1, then MAXTB, MAXLR must not be zero

Appendix A: DEFAULT TEST SETUP

Except where otherwise specified, all tests will require the DUT to have one of the following default physical configuration at the beginning of each test case:

Test Setup for NVMe Device:



Appendix B: NOTES ON TEST PROCEDURES

There are scenarios where in test procedures it is desirable to leave certain aspects of the testing procedure as general as possible. In these cases, the steps in the described test procedure may use placeholder values, or may intentionally use non-specific terminology, and the final determination of interpretation or choice of values is left to the discretion of the test technician. The following is an attempt to capture and describe all such instances used throughout the procedures.

<p>Ports on Testing Station and Device Under Test</p>	<p>In general, any PCIe Port on the Testing Station or Device Under Test may be used as an interface with a test station or interoperability partner. There is <i>assumed</i> to be no difference in behavior, with respect to the protocols involved in this test suite, between any two PCIe ports on the Testing Station or Device Under Test. Hence, actual ports used may be chosen for convenience. However, it is recommended that the PCIe port used in the test configuration is recorded by the test technician.</p>
<p>Use of “various”</p>	<p>To maintain generality, some steps will specify that “various other values” (or the like) should be used in place of a given parameter. Ideally, all possible values would be tested in this case. However, limits on available time may constrain the ability of the test technician to attempt this. Given this, a subset of the set of applicable values must generally be used.</p> <p>When deciding how many values should be used, it should be noted that the more values that are tested, the greater the confidence of the results obtained (although there is a diminishing return on this).</p> <p>When deciding which specific values to use, it is generally recommended to choose them at pseudo-randomly yet deterministically. However, if there exist subsets of the applicable values with special significance, values from each subset should be attempted.</p>

Appendix C: TEST TOOLS

The Tests described in this document can be performed using available IOL INTERACT NVMe Test Software available from UNH-IOL.

If using the PC Edition of the IOL INTERACT NVMe Test Software, UNH-IOL recommends using v16.0 or higher of the IOL INTERACT NVMe Test Software. This software is available via <https://www.iol.unh.edu/solutions/test-tools/interact>.

If using the Teledyne-LeCroy edition of the IOL INTERACT NVMe Test Software, UNH-IOL recommends using v16.0 or higher of the IOL INTERACT NVMe Test Software. This software is available at <https://www.iol.unh.edu/solutions/test-tools/interact>. This should be used in conjunction with the latest Teledyne-LeCroy software available at: <http://teledynelecroy.com>

Appendix D: NVME INTEGRATORS LIST REQUIREMENTS

Purpose: To provide guidance on what tests are required for NVMe Integrators List Qualification

References:

[1] NVMe Integrators List Policy Document

Resource Requirements:

NVMe Host Platform and Device.

Last Modification: April 2, 2019

Discussion: Each Test defined in this document is defined as being Mandatory (M), FYI, or In Progress (IP). This primary designation is shown in the title of the test case and is understood to apply to PCIe based products. An additional designation is provided if a test is applicable to NVMe-oF products (OF). Tests that are designated as being applicable to NVMe-oF Products are understood to inherit the primary designation of the test (i.e. M, FYI, IP), unless an additional designation is specified. The following examples are provided:

Test 1.1 Example Name (M)– Test is mandatory for all PCIe based products and does not apply to NVMe-oF products.

Test 2.1 Example Name (M, OF)– Test is mandatory for all products, including NVMe-oF products.

Test 3.1 Example Name (M, OF-IP)- Test is mandatory for all PCIe based products, and test is currently On Progress for NVMe-oF products.

NVMe protocol testing is independent of the transport used. Conformance tests described in this document may be performed at any link speed, width, or transport type that the NVMe product under test supports.

If a Test is designated as Mandatory, a product must pass this test in order to qualify for the NVMe Integrators List. For tests that deal with features defined as optional in the NVMe specification, a check is performed at the beginning of the test to determine if the optional feature is supported or not. If the optional feature is not supported the test is marked as ‘Not Applicable’ and does not impact qualification for the Integrators List.

If a Test is designated as FYI, a device does not need to pass this test in order to qualify for the NVMe Integrators List. Tests designated as FYI may become Mandatory tests in the future.

If a Test is designated as In Progress, a device does not need to pass this test in order to qualify for the NVMe Integrators List. These test cases are still under development. Tests designated as In Progress may become Mandatory tests in the future.

Any Test may have a Case within it with a different designation as the Test itself (i.e. a Mandatory test may include FYI cases). In this case, only the Mandatory Cases are required for NVMe Integrators List qualification.