

UNH-IOL NVMe Testing Service

Test Suite for NVMeoF Interoperability
Version 10.0



Last Updated: August 30, 2018

*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

MODIFICATION RECORD	3
ACKNOWLEDGMENTS	4
INTRODUCTION	5
REFERENCES	7
GROUP 1: FABRIC INTEROP TESTS (MULTI SWITCH).....	8
Test 1.1 – NVMeoF Initiator to Target verification in a Fabric (multi-switch).....	9
Test 1.2 – NVMeoF Initiator to Target traffic testing (multi-switch).....	10
Test 1.3 – NVMeoF Initiator to Target traffic Link Pull in Fabric (multi-switch)	11
Test 1.4 – NVMeoF Initiator to Target traffic from power up state in a Fabric (multi-switch)	13
Appendix A – Test Topologies	15
Appendix B – UNH-IOL Interop Test Bed.....	15
Appendix C – NVMe Fabrics Integrators List Requirements (EXAMPLE OF NVMe list requirements)	17
Appendix D – Other potential test cases	18

MODIFICATION RECORD

2017 June 14 (Version 7.9)

Tim Sheehan : Initial draft

2017 August 29 (Version 8.0)

Tim Sheehan : Initial Release

2017 September 12 (Version 8.0a)

David Woolf : Added FC and Ethernet Switches as eligible products in Appendix C.

2018 August 30 (Version 10.0)

David Woolf : 10.0 Release.

ACKNOWLEDGMENTS

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

Tim Sheehan

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 NVMeoF functionality of their products. This test suite is aimed at validating products in support of the work being directed by the NVMe Organization and Fibre Channel T11 and FCIA.

These tests are designed to determine if a product interoperates with other products designed to the Fibre Channel standards, hereafter referred to as the “NVMeoF Specifications” and NVM Express Revision 1.2 specification, hereafter referred to as the “NVMe Specification”). This interoperability of NVMe products happens across a data center fabric whose transport is the Fibre Channel protocol. Successful completion of these tests provide a reasonable level of confidence that the Device Under Test (DUT) will function properly in many NVMe environments.

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 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:

1. NVM Express Revision 1.3 (May 1, 2017)
2. NVM Express over Fabrics Revision 1.0 (June 5, 2016)
3. FC-NVMe INCITS 540-2018

Group 1: Fabric Interop Tests (Multi Switch)

Overview: This section describes methods for performing interoperability verification for NVMe products using a multi-switch Fabric topology between initiators and targets.

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 – NVMeoF Initiator to Target verification in a Fabric (multi-switch)

Purpose: To verify that an NVMe Initiator can properly discover a NVMe Storage Device in a Fibre Channel or Ethernet Fabric.

References:

- [1] NVMeoF Specification
- [2] FC Specifications

Resource Requirements:

- NVMeoF Host Platform and Device.
- 2 or more switches (Ethernet or FC)

Last Modification: August 15, 2017

Discussion: An NVMeoF Initiator should be able to discover a NVMeoF Target in a multi switch fabric topology and display this information in the Initiator operating system.

Test Setup: See Appendix A.

Test Procedure:

1. Build multi switch fabric.
2. Connect NVMeoF Initiator and Target via a multi switch fabric.
3. Allow NVMeoF Initiator to discover all fabric attached targets.

Observable Results:

1. Verify that the NVMeoF Target is visible in the Host Operating System of the NVMeoF Initiator.

Possible Problems: None known.

Test 1.2 – NVMeoF Initiator to Target traffic testing (multi-switch)

Purpose: To verify that an NVMe Initiator and Target in a Fabric (multi-switch) can transfer data without errors.

References:

- [1] NVMeoF Specification
- [2] FC Specifications

Resource Requirements:

- NVMeoF Host Platform and Device.
- 2 or more switches (Ethernet or FC)

Last Modification: August 15, 2017

Discussion: A NVMeoF Initiator and Target that are connected successfully at the Fabric and NVMe levels in a multi-switch fabric should be able to perform basic data traffic read, writes, read/writes.

Test Setup: See Appendix A. All NVMeoF Initiators connected to the fabric have discovered all available NVMeoF Targets.

Test Procedure:

1. Perform Write IO from the Initiator to the Target for 15 minutes. All blocks written should be Read back by the Initiator and Compared for data integrity.

Observable Results:

1. Ensure that no data integrity errors were detected during the IO operations.

Possible Problems: possible

Test 1.3 – NVMeoF Initiator to Target traffic Link Pull in Fabric (multi-switch)

Purpose: To verify that an NVMe Initiator can properly discover an NVMe Storage Device after link pulls in a fabric topology.

References:

- [1] NVMeoF Specification
- [2] FC Specifications

Resource Requirements:

- NVMeoF Host Platform and Device.
- 2 or more switches (Ethernet or FC)

Last Modification: August 15, 2017

Discussion: A NVMeoF Initiator and Target that are connected in a multi switch fabric topology can successfully discover NVMe targets and establish NVMe connectivity, should be able to perform basic data traffic read, writes, read/writes after the link between the Initiator and Target is removed and re-established.

Test Setup: See Appendix A. All NVMeoF Initiators connected to the fabric have discovered all available NVMeoF Targets.

Test Procedure:

1. Remove link from the NVMeoF Initiator by disconnecting the cable or optical module at the NVMeoF Initiator port.
2. Leave the link disconnected until the NVMeoF Host Operating System indicates that the NVMeoF target is no longer visible.
3. Reconnect the cable or optical module to the NVMeoF Initiator
4. Allow NVMeoF Initiator to discover all fabric attached targets.
5. When the NVMeoF Target is again visible from the NVMeoF Host Operating System, begin IO from the NVMeoF Initiator to the NVMeoF Target, let IO operations continue for no more than one minute.
6. Ensure no data integrity errors are detected.
7. Remove link from the NVMeoF Target by disconnecting the cable or optical module at the NVMeoF Target port.
8. Leave the link disconnected until the NVMeoF Host Operating System indicates that the NVMeoF target is no longer visible.
9. Reconnect the cable or optical module to the NVMeoF Target.
10. Allow NVMeoF Initiator to discover all fabric attached targets.
11. When the NVMeoF Target is again visible from the NVMeoF Host Operating System, begin IO from the NVMeoF Initiator to the NVMeoF Target, let IO operations continue for no more than one minute.
12. Ensure no data integrity errors are detected.
13. Remove link between the switches by disconnecting the cable or optical module at a switch port.
14. Leave the link disconnected until the NVMeoF Host Operating System indicates that the NVMeoF target is no longer visible.
15. Reconnect the cable or optical module to the switch.
16. Allow NVMeoF Initiator to discover all fabric attached targets.
17. When the NVMeoF Target is again visible from the NVMeoF Host Operating System, begin IO from the NVMeoF Initiator to the NVMeoF Target, let IO operations continue for no more than one minute.
18. Ensure no data integrity errors are detected.

Observable Results:

1. Verify that in each link removal instance, the NVMeoF Target was not visible from the NVMeoF Host Operating system after the link was removed.
2. Verify that in each link reconnection instance, the NVMeoF Target was visible from the NVMeoF Host Operating system after the link was reconnected.

3. Verify that no data integrity errors were reported during IO operations.

Possible Problems: None known.

Test 1.4 – NVMeoF Initiator to Target traffic from power up state in a Fabric (multi-switch)

Purpose: To verify that an NVMe Initiator can properly identify an attached NVMe Storage Device from a power up state of all elements in the fabric.

References:

- [1] NVMeoF Specification
- [2] FC Standards

Resource Requirements:

- NVMeoF Host Platform and Device.
- 2 or more switches (Ethernet or FC)

Last Modification: August 15, 2017

Discussion: A NVMeoF Initiator and Target that are connected successfully at the Fabric and NVMe levels in a multi-switch fabric topology should be able to perform basic data traffic read, writes, read/writes after the link between the Initiator and Target and switch are established after a successful power up of all three devices.

Test Setup: See Appendix A. All NVMeoF Initiators connected to the fabric have discovered all available NVMeoF Targets.

Test Procedure:

1. Power off the NVMeoF Initiator.
2. Power on the NVMeoF Initiator
3. Allow NVMeoF Initiator to rediscover all fabric attached targets.
4. When the NVMeoF Target is again visible from the NVMeoF Host Operating System, begin IO from the NVMeoF Initiator to the NVMeoF Target, let IO operations continue for no more than one minute.
5. Ensure no data integrity errors are detected.
6. Power off the NVMeoF Target.
7. Leave the NVMeoF Target powered off until the NVMeoF Host Operating System indicates that the NVMeoF target is no longer visible.
8. Power on the NVMeoF Target.
9. Allow NVMeoF Initiator to discover all fabric attached targets.
10. When the NVMeoF Target is again visible from the NVMeoF Host Operating System, begin IO from the NVMeoF Initiator to the NVMeoF Target, let IO operations continue for no more than one minute.
11. Ensure no data integrity errors are detected.
12. Power off Switch 1.
13. Leave the switch powered off until the NVMeoF Host Operating System indicates that the NVMeoF target is no longer visible.
14. Power on Switch 1.
15. Allow NVMeoF Initiator to discover all fabric attached targets.
16. When the NVMeoF Target is again visible from the NVMeoF Host Operating System, begin IO from the NVMeoF Initiator to the NVMeoF Target, let IO operations continue for no more than one minute.
17. Ensure no data integrity errors are detected.
18. Power off Switch 2.
19. Leave the switch powered off until the NVMeoF Host Operating System indicates that the NVMeoF target is no longer visible.
20. Power on Switch 2.
21. Allow NVMeoF Initiator to discover all fabric attached targets.
22. When the NVMeoF Target is again visible from the NVMeoF Host Operating System, begin IO from the NVMeoF Initiator to the NVMeoF Target, let IO operations continue for no more than one minute.
23. Ensure no data integrity errors are detected.


Observable Results:

1. Verify that in each power off instance, the NVMeoF Target was not visible from the NVMeoF Host Operating system after the power was removed.
2. Verify that in each power on instance, the NVMeoF Target was visible from the NVMeoF Host Operating system after power was restored.
3. Verify that no data integrity errors were reported during IO operations.

Possible Problems: None known.

Appendix A – Test Topologies

1. Multiple Switch Fabric



Appendix B – UNH-IOL Interop Test Bed

Purpose: To provide guidance on products that may be potential interop partners.

References:

[1] NVMeoF Specification

Resource Requirements:

NVMeoF Host Platform and Device.

Last Modification:

Discussion: NVMe Interoperability should be performed across a variety of host system hardware and operating systems. UNH-IOL maintains a test bed of potential interop partners for use during private in-lab testing and plugfest events. Currently products available in the interop test bed can be seen here:

- <https://www.iol.unh.edu/testing/storage/nvme/equipment>
- <https://www.iol.unh.edu/testing/switching/esp/equipment>
- <https://www.iol.unh.edu/testing/storage/fc/equipment>

Appendix C – NVMe Fabrics 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:

Discussion: The following table summarizes the interop test requirements, showing which items are optional or mandatory, as well as how many interop partners a test must be performed with. If a Product Under Test indicates support for a feature tested by an Optional test, that test must be performed and the Product Under Test must pass that test if the feature is supported.

Test Name	Mandatory or Optional	Required number of Interop Partners
1.1 – Target Verification	Mandatory	Pass with 2 or more Interop Partners
1.2 – Traffic Testing	Mandatory	Pass with 2 or more Interop Partners
1.3 – Link Pull in Fabric	Mandatory	Pass with 2 or more Interop Partners
1.4 – Power Up State	Mandatory	Pass with 2 or more Interop Partners

If the product under test is a target, among the 2 or more passing Interop Partners, there must be more than 1 HBA or OS represented.

If the product under test is an NVMeoF Initiator (ex RoCE, FC HBA) or an NVMeoF Switch (ex Ethernet Switch, Fibre Channel Switch, etc...), among the 2 or more passing Interop Partners, there must be more than one Target represented.

Appendix D – Other potential test cases

Overview: This section describes methods for performing interoperability verification for NVMe products using a Fibre Channel Fabric topology. These tests are added here to make people aware of them and to gather consensus if they are pertinent to interoperability testing.

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).

- 1. Buffer to Buffer credit recovery(CRC error injection)**
- 2. Multi-pathing**
- 3. Gen5 & 6 interoperability (Speed)**
- 4. Performance (SCSI .vs. NVMe)**
- 5. Combination of FC & NVMeoF**
- 6. NVMe Queue depth testing**
- 7. Switch trunking**
- 8. High availability**
- 9. Fabric FDMI-2 validation**
- 10. Multi-target environment**