



University of New Hampshire
InterOperability
Laboratory

The NVMe and NVMe-oF Plugfest #9 Webinar

Presented by: Kerry Munson

www.iol.unh.edu

UNH-IOL



Kerry Munson
Operations Manager,
Datacenter Technologies



David Woolf
Senior Engineer,
Datacenter Technologies

Who We Are



- Founded in 1866
- Main campus located in Durham, New Hampshire, USA
- 12,500+ undergraduate students, 2,500+ graduate students

UNH InterOperability Lab

- The UNH-IOL is a non-profit neutral, third-party laboratory dedicated to testing data networking technologies through industry collaboration.
- Since 2012, NVMe Org has collaborated with UNH-IOL to manage the Integrator's List based Test Program on behalf of NVMe Organization



University of New Hampshire
InterOperability
Laboratory

NVMe Plugfests

NVMe Plugfest #9

- June 11-14, 2018 @ UNH-IOL Durham, NH
- Registration closes May 21, 2018

NVMe-oF Plugfest

- June 18-21, 2018 @ UNH-IOL Durham, NH
- Registration closes May 29, 2018

NVMe Plugfest

- July 23-26, 2018 HUST Wuhan, China
- Registration closes May 25, 2018
- Only held if sufficient interest




Tests Available




NVMe
Conformance
Test Cases

220




NVMe
Interop
Test Cases

9




NVMe-MI
Conformance
Test Cases

53



NVMe-oF
Conformance
Test Cases

132



NVMe-oF
Interop
Test Cases

4

Glossary


M – Mandatory. Each test has a check to see if an option is supported or not

FYI –Not mandatory at this time

IP – In progress. Test still under development.


OF – Test available for NVMe-oF products

RC – Root Complex




NVMe

Pass all 'M' Conformance Test Cases
Pass IO Interop w/ 5 RCs
Pass Boot Interop w/2 RCs
Pass Hotplug w/1 RC



NVMe-MI

Pass all 'M' Conformance Test Cases



NVMe-oF

Pass all 'OF' Conformance Test Cases
Pass Interop with 2 partners

Agenda

- **IOL INTERACT Test Tool Updates**
- Test updates for:
 - NVMe PCIe SSDs for v1.3 for Conformance and Interop
 - NVMe-MI
 - NVMe over Fabrics
- Logistics
- Q&A

Conformance Tools



University of New Hampshire
InterOperability
Laboratory

Conformance Test Suite, Interop Test Suite, Plugfest, Integrators List Version, Software tools move forward together on Major Release, roughly twice a year: 7>8 (Oct 2018), 8>9 (June 2018), 9>10 (2H18) etc...

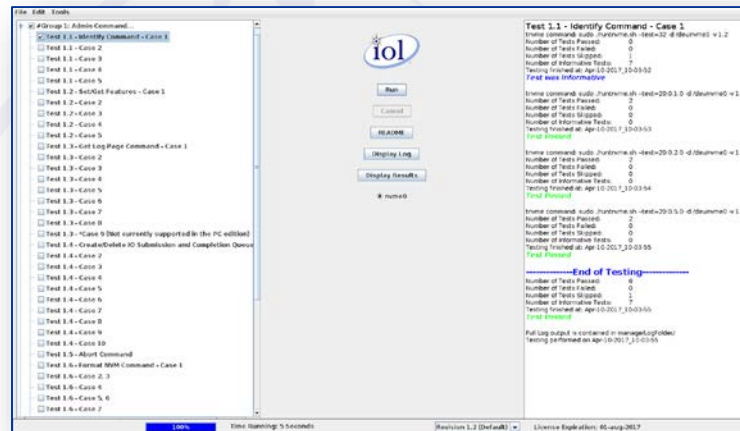
- New mandatory test requirements will only be introduced in Major Release, after vetting through TP process

Minor releases, which affect tools only, happen on 3-4 week intervals within Major Release, denoted by a minor letter (9.0a , 9.0b, 9.0c)

- New FYI test functionality may be introduced
- New usability features
- Mandatory tests unchanged aside from Bug fixes
- Tool updates may not correspond to one another

UNH-IOL NVMe INTERACT Test Tools

- PC Edition for PCIe, RoCE, FC released
- Teledyne-LeCroy Edition for PCIe released
- Version 9.0a or higher will be used at Plugfest #9
 - Version 8.0 available
- Download today from UNH-IOL Box Fileshare and start testing




Conformance Tools



University of New Hampshire
InterOperability
Laboratory

NVMe Conformance

- UNH-IOL PC Edition – PCIe 9.0a+
- Some tests also available in UNH-IOL Teledyne-LeCroy Edition




NVMe-MI Conformance

- Today uses Teledyne-LeCroy T34 and PETracer SW 8.76



NVMe-oF Conformance

- UNH-IOL PC Edition - RoCE 9.0a+
- UNH-IOL PC Edition – FC 9.0a+



Bug Tracking Improvements



University of New Hampshire
InterOperability
Laboratory

- Bugs can be reported through the UNH-IOL website
- When a bug is reported, reporter will receive a bug number, and estimated fix date based on next release schedule.
- The release notes will show what bugs (with their bug numbers) have been addressed in that release
- Enable user to make a connection to bug reported > bug fixed
- Avoid “engineering release”, except in extreme cases

IOL INTERACT PC Edition PCIe Package 9.0a	
Issue Key	Description
NVMELDD-49	Modified Test 4.18 to check that VS Structure matches VS CAP
NVMELDD-50	Modified Test 2.1 case 4 to follow requirement for DUT to report low
NVMELDD-51	Modified Test 2.3 case 4 to follow requirement for DUT to report low
NVMELDD-52	Modified Test t 2.4 case 4 to follow requirement for DUT to report low
NVMELDD-53	Modified Test 2.3 case 7 to follow requirement for DUT to report low
NVMELDD-54	Modified Test 2.7 case 3 to follow requirement for DUT to report low
NVMELDD-55	Modified Test 2.7 case 5 to verify c

Agenda

- IOL INTERACT Test Tool Updates
- Test updates for:
 - **NVMe PCIe SSDs for v1.3 for Conformance and Interop**
 - NVMe-MI
 - NVMe over Fabrics
- Logistics
- Q&A

NVMe Conformance Test Requirements

- Test both PC Edition and LeCroy Edition IOL INTERACT conformance tests
- If your product passes on both PC edition and LeCroy edition, *no further review is required*
- PC edition and LeCroy edition conformance testing each take approximately 30 minutes
- The latest v9.0 test suites, along with a redline version are available at: <https://www.iol.unh.edu/testing/storage/nvme/test-suites>

NVMe Conformance Test Updates

NVMe FYI > Mandatory changes

- **Test 1.4 Create/Delete I/O Submission and Completion Queue Commands Case 10: Create I/O Submission Queue Invalid Queue Address Offset** recommended to be Mandatory for NVMe Drives.
- **Test 1.7 Asynchronous Events Case 1: Asynchronous Event Request Command** recommended to be Mandatory for NVMe Drives.
- **Test 2.2 Dataset Management Command - Case 3: Deallocate Out of Range, Case 4: NR Value is Maximum, and Case 5: Correct Range Deallocated** are recommended to be Mandatory for NVMe Drives.

All Tests proposed to go FYI>Mandatory were run successfully at Plugfest #8



NVMe Conformance Test Updates



NVMe FYI > Mandatory changes

- **Test 2.3 Read Command Case 11: Valid READ, Invalid PRP Address Offset** recommended to be Mandatory for NVMe Drives.
- **Test 2.4 Write Command Case 11: Write with Invalid PRP Address Offset** recommended to be Mandatory for NVMe Drives.
- **Test 2.7 Write Zeroes Command Case 9: PRCHK is Non Zero** recommended to be Mandatory for NVMe Drives.
- **Test 2.9 AWUN/NAWUN Case 2: Atomic Boundaries Supported** recommended to be Mandatory for NVMe Drives.
- **Test 5.5 Command Specific Errors Definition Case 6: Feature Not Namespace Specific IV=1** recommended to be Mandatory for NVMe Drives.

NVMe Conformance Test Updates



New NVMe Conformance Test Cases

- 5 new cases for Telemetry Log Pages added under *Test 1.3 Get Log Page*
- 8 new cases for Device Self Test Short Operations added under new *Test 1.10 Device Self Test Short Operations*
- 8 new cases for Device Self Test Extended Operations added under new *Test 1.11 Device Self Test Extended Operations*
- 5 new cases for Aborted Device Self Test Short Operations added under new *Test 1.12 Aborted Device Self Test Short Operations*
- 4 new cases for Aborted Device Self Test Extended Operations added under new *Test 1.13 Aborted Device Self Test Extended Operations*

New Test
Cases always
FYI

DST support is spotty
Devices return completions
immediately, clear that no DST is
being performed

Samples
Samples
Samples

NVMe Conformance Test Updates



New NVMe Conformance Test Cases

- 1 new case for Namespace Identification Descriptors added under *Test 1.1 Identify*
- Added Timestamp Feature to features checked in *Test 1.2 Set/Get Feature*
- 1 new cases for Non-operation state Admin Commands added under *Test 8.2 Return from Non-Operation State*
- 1 new cases added for *Test 8.4 Power State Entrance Latency*
- 1 new cases added for *Test 8.5 Power State Exit Latency*
- 1 new cases added for *Test 8.6 Relative Read Throughput*
- 1 new cases added for *Test 8.7 Relative Write Throughput*
- 2 new cases added for *Test 8.8 Host Controlled Thermal Management*, one for basic operation and one for Invalid Field detection.



New Test
Cases
always
FYI

NVMe Conformance Test Updates



Updated NVMe Conformance Test Cases

- Many clarifications and updates captured in Redline
- *2 primary update types:*
 1. Adjusted tests to not expect error codes to be reported in numerical order
 2. Adjusted tests to include appropriate checks of OACS, ONCS etc to check for feature support prior to test execution

NVMe Interop Test Requirements

- No significant changes to Interop Tests
 - Read/Write + Boot with 5 systems
 - Hotplug if U.2 with 1 system
- Interop Read/Write MLTT Scripts shared on UNH-IOL Box Fileshare
- Interop testing has been automated to help save time during interop test blocks
- Interop testing will take approximately 90 minutes
- The latest v9.0 test suite available at:
<https://www.iol.unh.edu/testing/storage/nvme/test-suites>



NVMe Interop Updates



NVMe Interop New Test

- Test 1.9 Return from Hibernation
- All new tests are FYI
- Test involves cycling the host through OS enabled hibernation states, and ensuring drive wakes up properly.

Interop Systems

Platform	OS	Testing
Dell PowerEdge R720	Win 2k12 Server R1 & R2	Read/write
Dell PowerEdge R720	RHEL 6.5	Hot plug
ASUS X79 Deluxe LGA 2011	Ubuntu 16.04	Read/write
SUPERMICRO MBD-X10SAT-O ATX Server	RHEL 6.5	Read/write
SUPERMICRO MBD-X10SAT-O ATX Server	Windows 10	Read/write
ASRock Z97	Ubuntu 16.04	Boot
Gigabyte GA-1 H170	Ubuntu 16.04	Boot
Intel 'Wildcat Pass'	Win2k16 Server	Read/write
Intel 'Wildcat Pass'	CentOS	Read/write
AMD Epyc	Ubuntu 16.04	Read/write
Dell PowerEdge R740XD Server	Windows 10	Read/write

Agenda

- IOL INTERACT Test Tool Updates
- Test updates for:
 - NVMe PCIe SSDs for v1.3 for Conformance and Interop
 - **NVMe-MI**
 - NVMe over Fabrics
- Logistics
- Q&A

NVMe-MI Conformance Test Requirements

- Test performed using the Teledyne-LeCroy T-34
 - Software 8.76
- Currently no required NVMe-MI Interop Tests
 - Proof of concept for NVMe-MI Interop
- The latest v9.0 test suite, is available at:
<https://www.iol.unh.edu/testing/storage/nvme/test-suites>
- **NVMe and NVMe-MI Integrator's Lists are Independent**



NVMe-MI Updates

NVMe-MI FYI > Mandatory changes

- **Test 4.10 – Invalid Command Size** recommended to be Mandatory
- **Test 5.3 – CRC Check** recommended to be Mandatory

Contact Teledyne-LeCroy for NVMe-MI tools

PoC on NVMe-MI Interop



Agenda

- IOL INTERACT Test Tool Updates
- Test updates for:
 - NVMe PCIe SSDs for v1.3 for Conformance and Interop
 - NVMe-MI
 - **NVMe over Fabrics**
- Logistics
- Q&A

NVMe-oF Integrator's List Requirements

NVMe-oF Integrator's List will accept:

- RoCE Initiators and Targets
- Ethernet Switches
- Fibre Channel Initiators and Targets, Switches

Requirements documented in [NVMe Integrators List Policy Document v9.0](#)

NVMe-oF Integrator's List Requirements

Requirements for Initiators:

- Pass NVMe-oF Interop Tests with 2 Target Interop Partners

Requirements for Switches:

- Pass NVMe-oF Interop Tests with 2 Target Interop Partners

Requirements for Targets:


- Pass NVMe-oF Interop Tests with 2 Initiator Interop Partners
- Pass NVMe Protocol Conformance 'OF' tests

2 Partners can be the same HW with different OS/Driver.

NVMe-oF Interop Tests

Pass NVMe-oF Interop Tests with 2 partners in a multi switch setup.

[NVMe-oF Interop Test Suite Document v9.0](#)

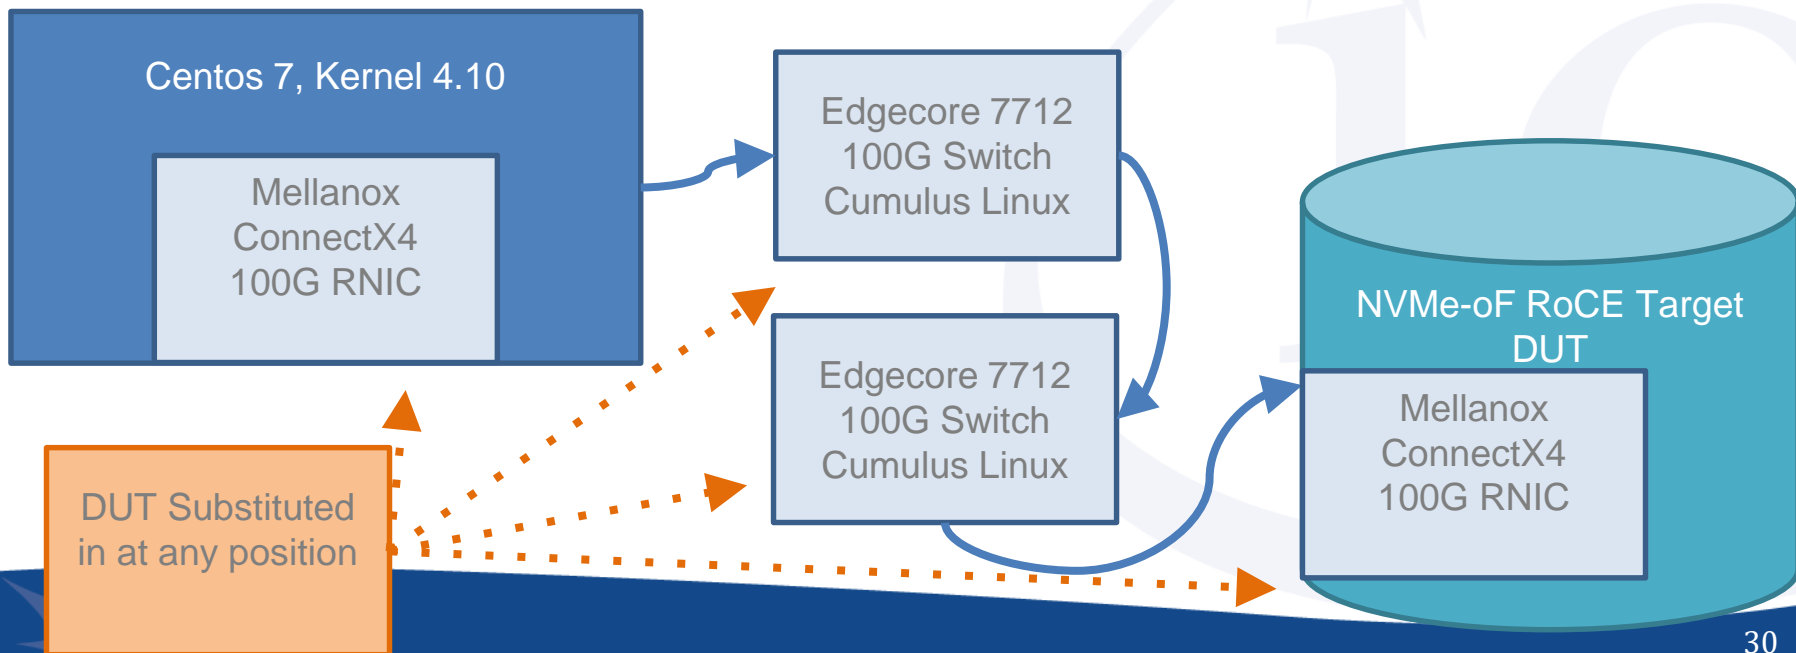


NVMe-oF Interop Tests Summary:

1. Discover Target
2. Send Traffic
3. Link Pull
4. Power Up

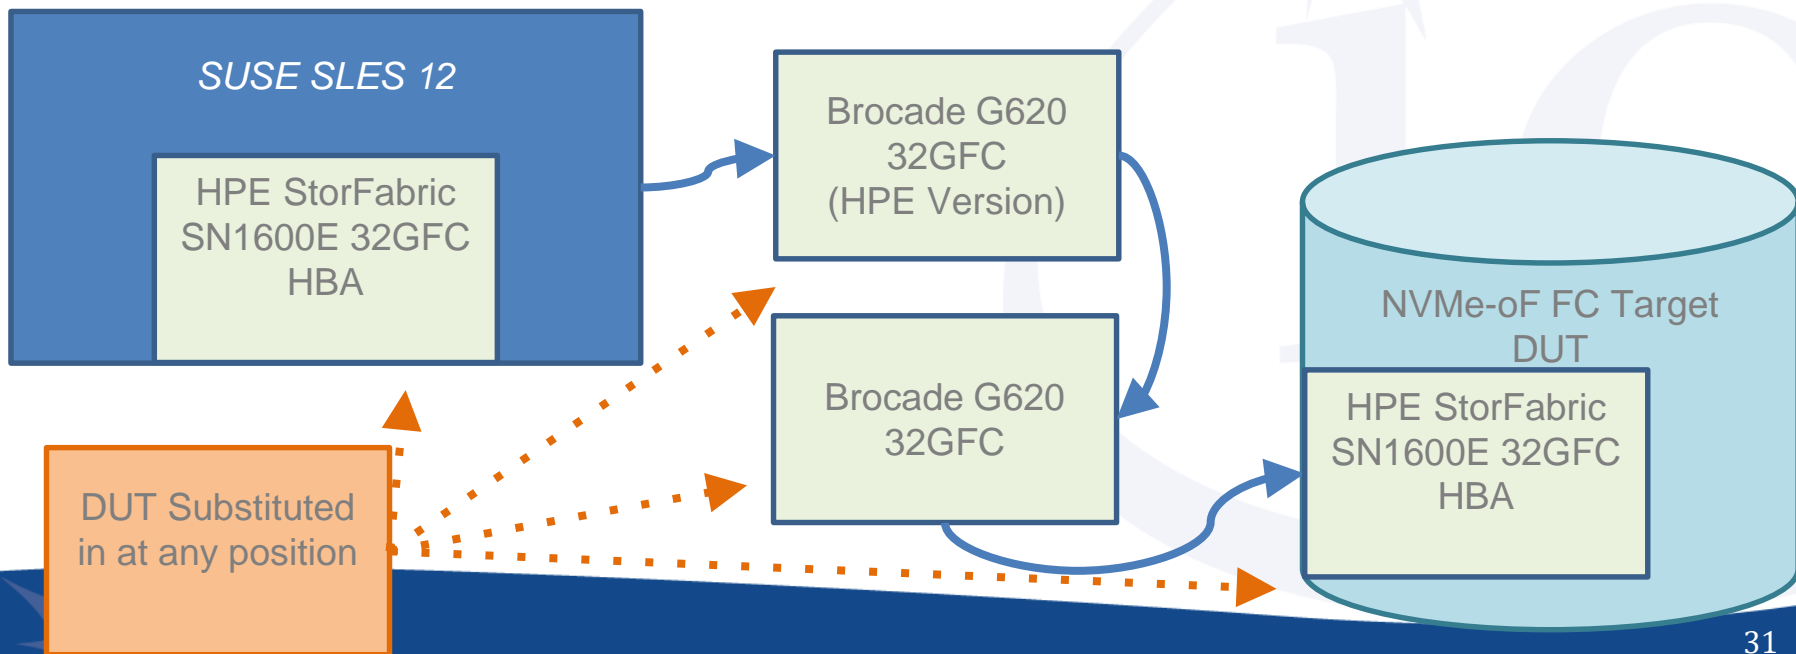
NVMe-oF Interop Tests

UNH-IOL NVMe-oF RoCE HW/SW Setup



NVMe-oF Interop Tests

UNH-IOL NVMe-oF Fibre Channel HW/SW Setup

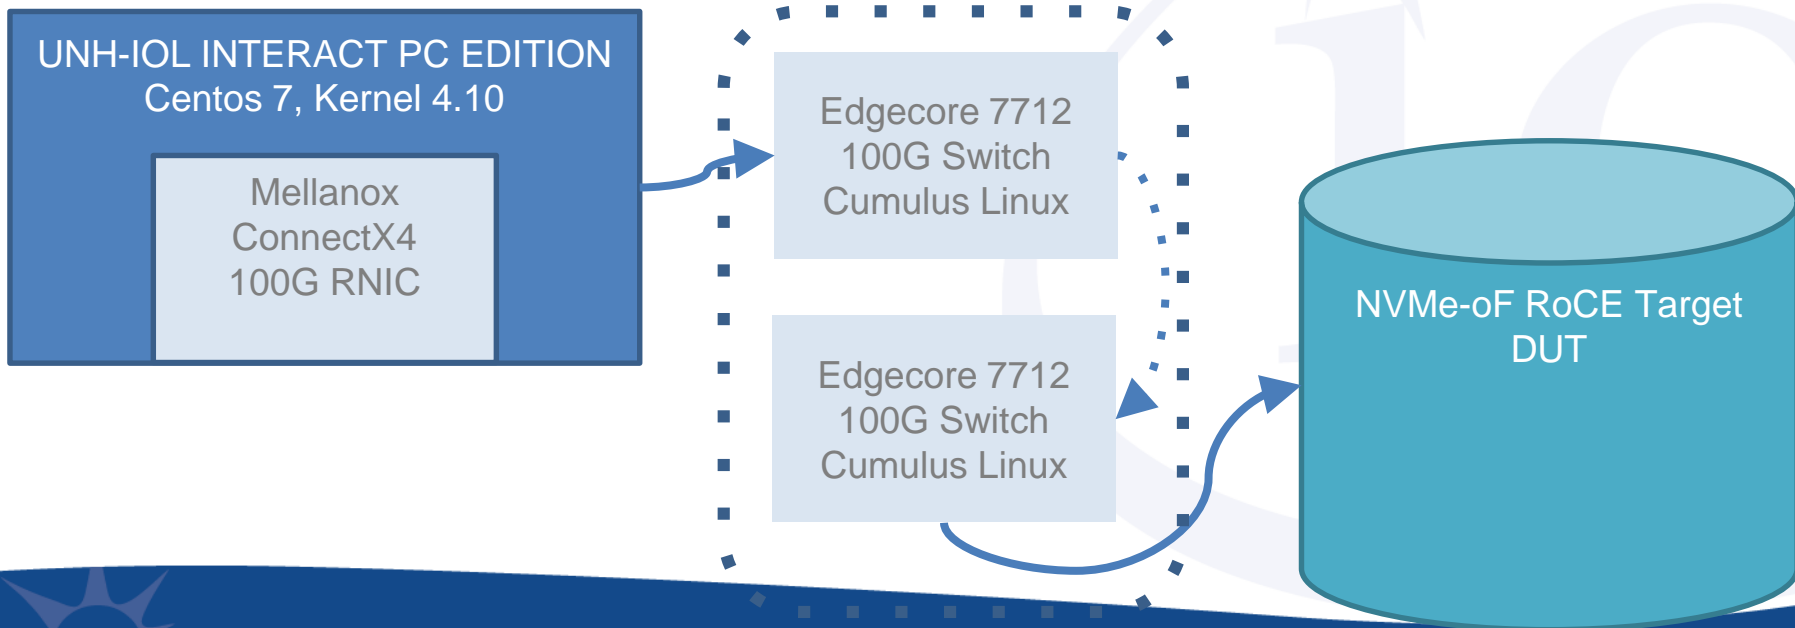


NVMe-oF: Conformance Tools

- UNH-IOL IOL INTERACT PC Edition for FC and RoCE
- Tool acts as NVMe-oF Initiator, to exercise the the NVMe protocol functionality on the Target.
- Encourage members to [download](#) and run tool ahead of event.

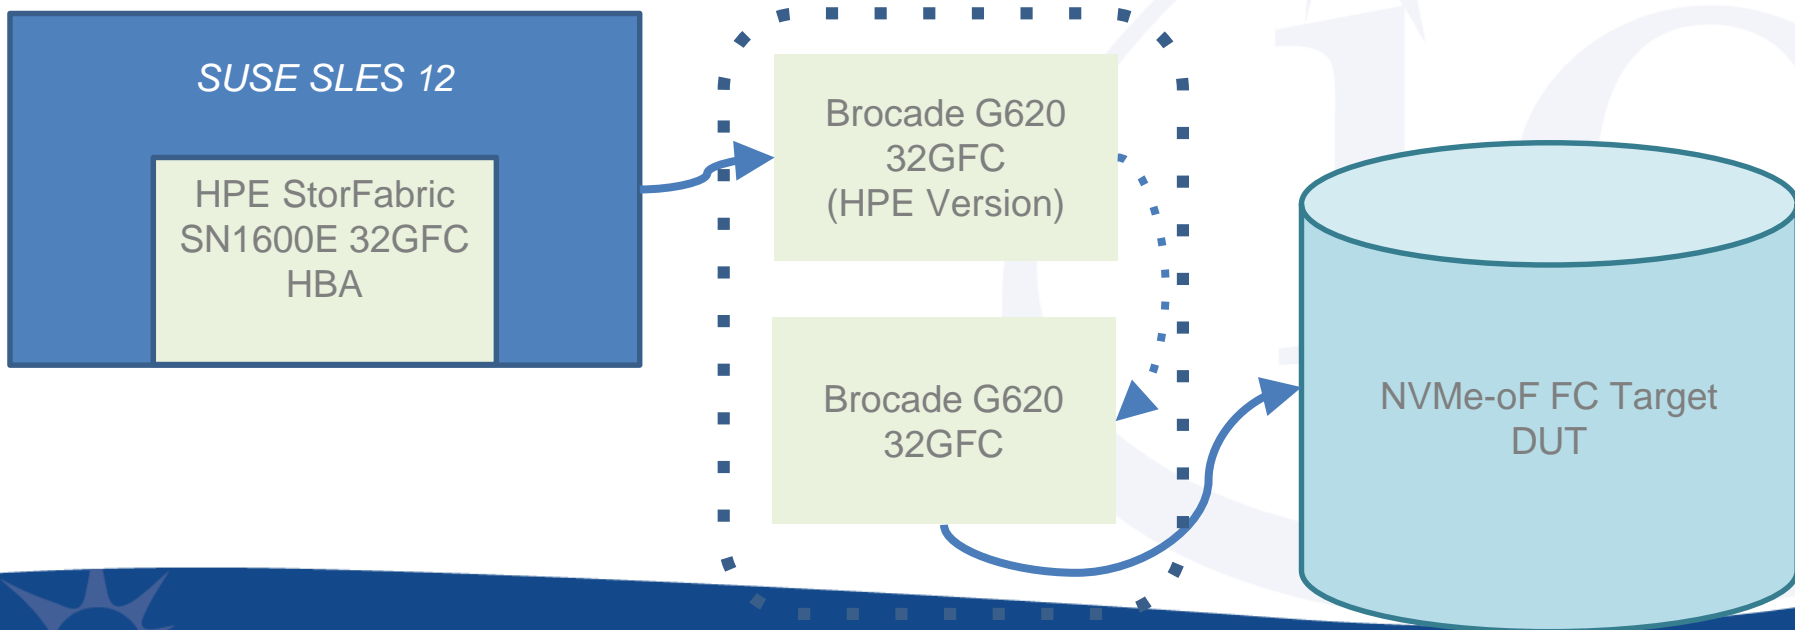
NVMe-oF: Conformance Tools

IOL INTERACT PC Edition RoCE HW/SW Setup



NVMe-oF: Conformance Tools

IOL INTERACT PC Edition FC HW/SW Setup



NVMe-oF Tests: NVMe Conformance Tests

- Documented in [NVMe Conformance Test Suite v9.0](#)
- Pass NVMe Protocol Conformance tests marked 'OF' (~60% of Tests)
- Currently verifies NVMe protocol only, no verification of Binding Spec

University of New Hampshire InterOperability Laboratory – NVMe O

Test 1.6 – Format NVM Command (MS, OF)

Purpose: To verify that an NVMe Controller can properly execute the Format N

References:

[1] NVMe Specification 5.15

Resource Requirements:

Tools capable of monitoring and decoding traffic on the NVMe interface

Last Modification: July 7, 2016

Discussion: The Format NVM command is used to low level format the NVM wants to change the LBA data size and/or metadata size. A low level format


NVMe-oF Updates



University of New Hampshire
InterOperability
Laboratory

NVMe-oF FYI > Mandatory changes

- **Test 1.1 Identify Command Case 1: Identify Namespace Data Structure and Case 2: Identify Controller Data Structure** recommended to be Mandatory for NVMe-oF NVM Subsystems.
- **Test 1.3 Get Log Page Command Case 1: Supported LIDs and Case 2: Unsupported Vendor Specific LIDs** recommended to be Mandatory for NVMe-oF NVM Subsystems.
- **Test 1.7 Asynchronous Events Case 2: Outstanding Commands Aborted after Reset** recommended to be Mandatory for NVMe-oF NVM Subsystems.
- **Test 2.5 Write Uncorrectable Command** all cases recommended to be Mandatory for NVMe-oF NVM Subsystems.



NVMe-oF Updates



University of New Hampshire
InterOperability
Laboratory

NVMe-oF FYI > Mandatory changes

- **Test 2.6 Flush Command Case 2: Invalid Namespace ID** recommended to be Mandatory for NVMe-oF NVM Subsystems.
- **Test 2.8 Atomicity Parameters** recommended to be Mandatory for NVMe-oF NVM Subsystems.
- **Test 3.3 Power Management Cases 1-4:** recommended to be Mandatory for NVMe-oF NVM Subsystems.
- **Test 5.1 Page Base Address and Offset (PBAO)** recommended to be Mandatory for NVMe-oF NVM Subsystems.

PoC on NVMe-oF Large Build

PoC on NVMe-oF / Infiniband (if sufficient interest)



Agenda

- IOL INTERACT Test Tool Updates
- Test updates for:
 - NVMe PCIe SSDs for v1.3 for Conformance and Interop
 - NVMe-MI
 - NVMe over Fabrics
- **Logistics**
- Q&A

Be Prepared!

Download the latest tools and run them!

- Tools that we will use at the plugfest are available now
- Run them, find bugs, fix bugs
- Link to licensing agreement
 - <https://www.iol.unh.edu/solutions/test-tools/interact>
- Link to UNH Box PC Edition IOL-INTERACT Software
 - <https://unh.box.com/s/cktgos25cjfk6alclromuvqfn1lgebqd>
- Link to UNH Box LeCroy Edition IOL-INTERACT Software
 - <https://unh.box.com/s/n3drj5qpmm7maqjrahxw3ycovuy3p37g>


Be Prepared!

Bring 2+ Samples of each product

- Conformance tests will only require 1 sample
- Interop tests require that **2 samples** be used simultaneously
- Please bring **at least 2 samples** of your product, this will facilitate interop testing
- Bring any necessary PCs, tools, cables that will be necessary to re-program or re-flash your device


Be Prepared!

- If you have a larger product (i.e. large PCIe dev board), please bring your own PCIe flex or riser cables



Be Prepared!

- UNH-IOL will have U.2/SFF-8639 to CEM adapters on hand
- If you have a **non CEM** form factor, please bring it, and also please bring at least 2 of your own adapters - this will facilitate interop testing
- Adapters available from *serialcables.com* and *teledynelecroy.com*



Available!

- We have a soldering station in the Plugfest Room.
- Wired and wireless internet provided
- Power cords, mice, keyboards, monitors



No Photos

- **No photos!**
- Please respect confidentiality of your colleagues.
- All companies have agreed to UNH-IOL Usage Agreement
- UNH-IOL may take some 'generic' photos.
- If you would like a photo of your team, or your own equipment setup, please check with UNH-IOL staff first.



Travel

Visa Invitation Letter Contact:

kerry.munson@iol.unh.edu

Shipping Address:

NVMe Plugfest % Kerry Munson
21 Madbury Rd Suite 100
Durham, NH 03824
☎: 603-862-0090

Equipment must arrive by
June 7th or June 14th

Airports: ✈

MHT - Manchester Boston Regional
Airport (🕒 1 hr)

BOS - Boston Logan International
Airport (🕒 1hr 30mins)

Train: 🚆

Amtrak - Downeaster
Boston North Station > Durham-UNH
<http://www.amtrakdowneaster.com/>

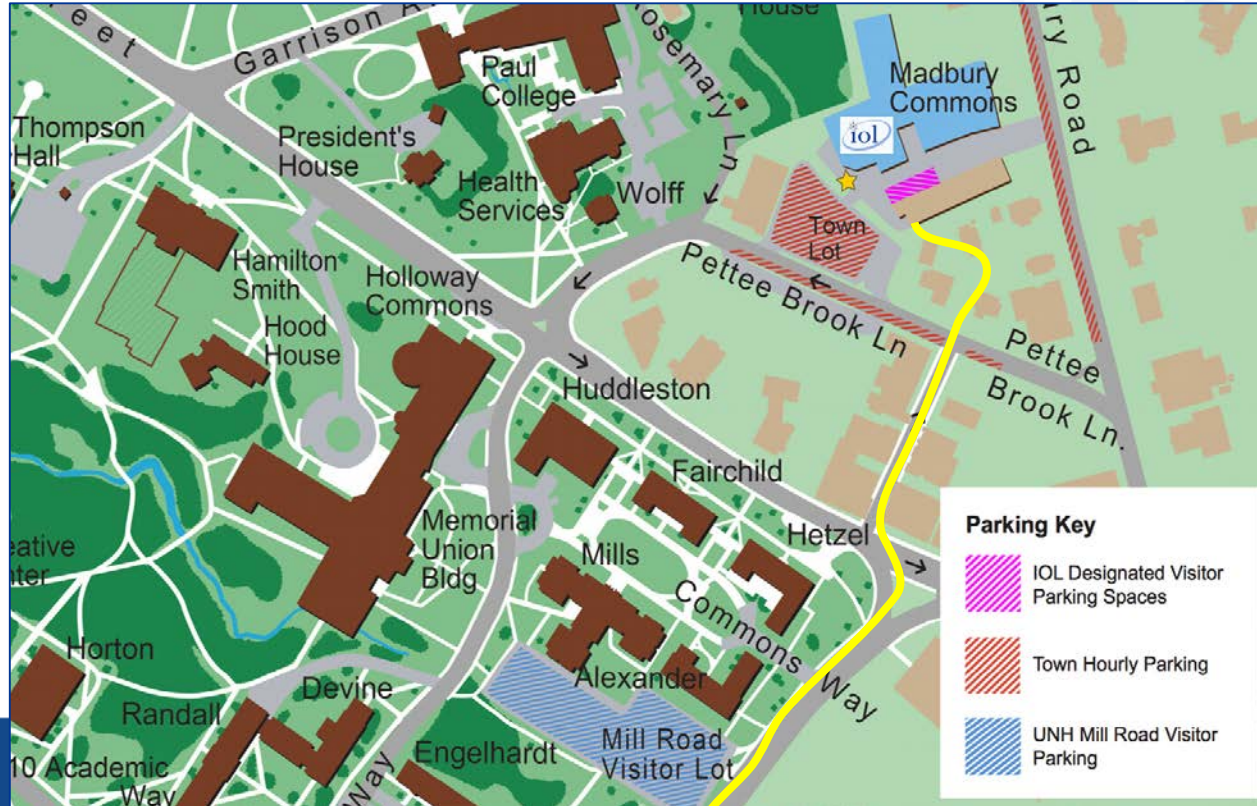
Hotels, Travel, Parking Info:

<https://www.iol.unh.edu/about/visit>

Parking

Parking code will be provided to registrants

[UNH Campus Map](#)



Event Logistics

- Plugfest runs Monday - Thursday
- Doors open at 8AM, testing until 6PM
- Kickoff meeting Monday @ 11AM
- Light breakfast and lunch provided each day
- Device schedule will be made after registration closes



Evening Reception!

Registration

Register here:

<https://www.iol.unh.edu/testing/storage/nvme/grouptest>

Registration will close May 21, 2018

Registration Limits:

- At registration 2 products may be registered.
- All other products will be waitlisted, and accepted on a space-available basis.
- At registration you can register 4 engineers.
- Additional engineers will incur an additional fee.

ICC Calls

Interop and Compliance Committee (ICC)


Chair: Ryan Holmqvist, Microsemi

Meets every other Tuesday 2PM Eastern / 11AM Pacific

Announcements/Meeting invites made on ICC reflector



University of New Hampshire
InterOperability
Laboratory



Email: Kerry.Munson@iol.unh.edu

Web: www.iol.unh.edu

Q&A

Q&A



University of New Hampshire
InterOperability
Laboratory

Q: Can you provide detailed info about the host platforms used in the plugfest tests for interop tests? for example, the motherboard make, model, BIOS, OS versions (Windows or Linux).

A: You can find this information at: <https://www.iol.unh.edu/testing/storage/nvme/equipment>

Q: Can members see the list of open bugs for the tools?

A: We do not have a way to share this right now.

Q: Regarding NVMe-MI conformance, are there any plans in the future to include other vendors for performing the conformance tests?

A: This is best addressed via the Interop and Compliance Committee.

Q: Can you add the link to the bug reporting tool?

A: You can report bugs in the tools via your MyIOL login: <https://www.iol.unh.edu/my/login>

Q&A



University of New Hampshire
InterOperability
Laboratory

Q: Is it necessary to configure end to end DCBx for NVMe-oF testing?

A: The integrators list tests do not cause any network congestions, so we won't require any DCBx. We can try turning on DCBx as an experiment for information gathering, so long as it doesn't interfere with Integrators List testing.

Q: Can we bring just our HBA card (PCie) and plug into your server or do we need to bring our server as well?

A: We would strongly recommend bringing your own server to house your HBA, to avoid any potential interop or configuration issues with using the HBA in an unfamiliar system at the plugfest.

Q: Will you be performing Host related testing ? Or only drive related?

A: The conformance testing is focused on PCIe SSDs and NVMe-oF Targets. The Interop Testing is designed for testing both Hosts and Targets.

Q&A



University of New Hampshire
InterOperability
Laboratory

Q: Did you try the PC tests on ARM + CentOS?

A: We have not tried this yet, but are eager to hear if anyone has tried it, and what their experience was. We'd be happy to add this to our supported systems for conformance testing.

Q: Regarding the PC Edition test tools, is this run on one particular system?

A: UNH-IOL will provide a system at the plugfest for running the PC Edition tests. Generally, we have not seen many problems running the tools on server or desktop systems based on regular Intel PCIe chipsets. However, more recently we have received some reports of problems when trying to use a Host which supports a lower number of Max IRQ than that of the drive being testing. Therefore, it is recommended to use a host system with a Max IRQ greater than or equal to the Max IRQ of the drive. If the host has a Max IRQ value less than that of the drive, This discrepancy can manifest itself in failing Test 5.5. Common chipsets like the Z97 chipset have a Max IRQ of 128, and some drives report Max higher IRQ values.

Q&A



University of New Hampshire
InterOperability
Laboratory

Q: Will there be a list of interested companies/parties for each section of testing?

A: We will provide a schedule of interop partners shortly after the registration for each event closes, so each participant will know who their interop partners will be well before the event.

Q: How is the automation helpful for interop testing?

A: We have automated some of the shutdown and reboot cycles during the interop test to avoid unnecessary user intervention and to save time

Q: How can we access the MLTT scripts?

A: We have shared these via our Box fileshare: <https://unh.app.box.com/folder/49286499020>