



OpenFabrics Alliance

Interoperability Logo Group (OFILG)

January 2016 Logo Event Report

UNH-IOL – 21 Madbury Rd., Suite 100 – Durham, NH 03824 – +1-603-862-0090
OpenFabrics Interoperability Logo Group (OFILG) – ofalab@iol.unh.edu

Ido Zemah
 Mellanox Technologies LTD
 Hermon Building 4th Floor
 P.O Box 586 Yokenam 20692
 Israel

Date: March 8, 2016
 Report Revision: 1.0
 OFED Version: 3.18-1
 OS Version: Scientific Linux 7.1

Enclosed are the results from OFA Logo testing performed on the following devices under test (DUTs):

Mellanox MCX314A-BCCT
Mellanox MCX312B-XCCT

The test suite referenced in this report is available at the UNH-IOL website. Release 2.04(2015-Jul -15) was used.

<https://iol.unh.edu/ofatestplan>

The following table highlights the Mandatory test results required for the OpenFabrics Interoperability Logo for the RoCE Network Adapter device class per the Test Plan referenced above and the current OpenFabrics Interoperability Logo Program (OFILP).

Test Procedures	IWG Test Status	Result/Notes
12.2: RoCE Link Initialization	Mandatory	Pass
12.4: IPoCE	Mandatory	Pass
13.4: TI uDAPL	Mandatory	Pass
13.5: TI RDMA Basic Interop	Mandatory	Pass
13.7: TI Rsockets	Mandatory	Pass
13.8: TI MPI – Open MPI	Mandatory	Pass

For specific details regarding issues, please see the corresponding test result.

Testing Completed March 8, 2016

Stefan Oesterreich
soesterreich@iol.unh.edu



Reviewed & Issued March 17, 2016

Bob Noseworthy
ren@iol.unh.edu

Result Summary

The Following table summarizes all results from the event pertinent to this RoCE device class (RCA).

Test Procedures	IWG Test Status	Result/Notes
12.2: RoCE Link Initialization	Mandatory	Pass
12.4: IPoCE	Mandatory	Pass
13.2: TI NFSoRDMA	Beta	Fail
13.4: TI uDAPL	Mandatory	Pass
13.5: TI RDMA Basic Interop	Mandatory	Pass
13.7: TI Rsockets	Mandatory	Pass
13.8: TI MPI – Open MPI	Mandatory	Pass

Digital Signature Information

This document was created using an Adobe digital signature. A digital signature helps to ensure the authenticity of the document, but only in this digital format. For information on how to verify this document's integrity proceed to the following site:

<http://www.iol.unh.edu/certifyDoc/>

If the document status still indicates "Validity of author NOT confirmed", then please contact the UNH-IOL to confirm the document's authenticity. To further validate the certificate integrity, Adobe 6.0 or later should report the following fingerprint information:

MD5 Fingerprint: 39 96 A4 8C 02 E0 38 DA F7 AB E1 ED BD E9 B5 F0

SHA-1 Fingerprint: 41 EE 65 F4 8A 6A 3A A6 3D AF 63 F4 78 7D 97 7B 48 49 C9 DD

Report Revision History

- v1.0 Initial Release

Configuration Files

Description	Attachment
Scientific Linux 7.1 Configuration File	
OFED 3.18-1 Configuration File	

Result Key

The following table contains possible results and their meanings:

Result:	Description:
PASS	The Device Under Test (DUT) was observed to exhibit conformant behavior.
PASS with Comments	The DUT was observed to exhibit conformant behavior however an additional explanation of the situation is included.
Qualified PASS	The DUT was observed to exhibit conformant behavior, with the exception of fault(s) or defect(s) which were previously known.
FAIL	The DUT was observed to exhibit non-conformant behavior.
Warning	The DUT was observed to exhibit behavior that is not recommended.
Informative	Results are for informative purposes only and are not judged on a pass or fail basis.
Refer to Comments	From the observations, a valid pass or fail could not be determined. An additional explanation of the situation is included.
Not Applicable	The DUT does not support the technology required to perform this test.
Not Available	Due to testing station limitations or time limitations, the tests could not be performed.
Borderline	The observed values of the specific parameters are valid at one extreme and invalid at the other.
Not Tested	Not tested due to the time constraints of the test period.

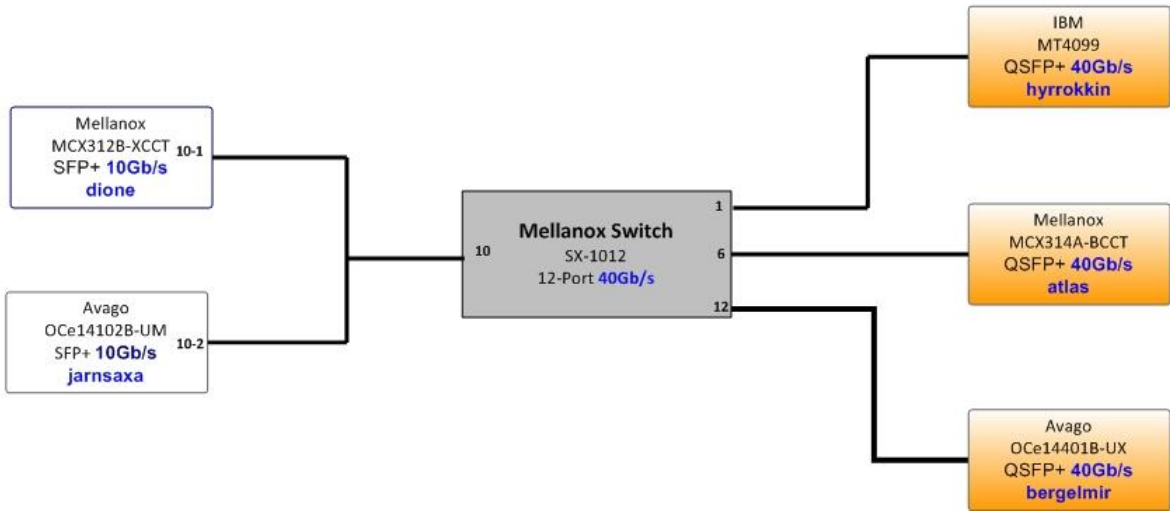
DUT and Test Setup Information

Figure 1: The RoCE fabric configuration utilized for all testing is shown below.

January 2016 Logo
 RoCE Topology

Ethernet Addressing
 <hostname>.ofa

RoCE Addressing
 <hostname>-ce.ofa



DUT #1 Details			
Manufacturer:	Mellanox	Firmware Revision:	2.32.5100
Model:	MCX314A-BCCT	Hardware Revision:	00
Speed:	40Gb/s	Located in Host:	atlas
Firmware MD5sum:	N/A		
Additional Comments / Notes:			

DUT #2 Details			
Manufacturer:	Mellanox	Firmware Revision:	2.32.5100
Model:	MCX312B-XCCT	Hardware Revision:	00
Speed:	10Gb/s	Located in Host:	dione
Firmware MD5sum:	N/A		
Additional Comments / Notes:			

Mandatory Tests – RoCE Device Test Results:

12.2: RoCE Link Initialization

Test Result	Pass
Result Discussion:	
All devices were seen to link at proper speeds. All ICMP requests were sent and received properly.	

12.4: IPoCE

Test Result	Pass
Result Discussion:	
All devices were seen to link at proper speeds. All ICMP requests were sent and received properly.	

13.2: TI NFSoRDMA

Test Result	Fail
Discussion:	
When the Power8 machine connects to 40G Mellanox Adapter as a client, the Power8 machine crashes when running the “mount” command. When the Power8 Machine acts as server and 40G Mellanox Adapter is the client, tests run successfully. It is unknown whether the fault lies on the Power8 machine or this DUT. When run against a second 10G Mellanox Adapter, all tests run successfully.	

13.4: TI uDAPL

Test Result	Pass
Discussion:	
With the use of the daplttest utility, devices exhibited simple send receive, verification, polling, scatter gather on a point-to-point topology. Devices were then connected to a switched topology and, with the use of the daplttest utility, tested on correct use of multiple threads, RDMA read and write operations, pipeline tests with RDMA write scatter gather list and RDMA read, and the multiple switches test.	

13.5: TI RDMA Basic Interoperability

Test Result	Pass
Discussion:	
All devices were shown to correctly exchange core RDMA operations across a simple network path under nominal (unstressed) conditions. Each HCA acted as both a client and a server for all tests listed below.	
<ul style="list-style-type: none">• Small & Large RDMA Read• Small & Large RDMA Write• Small & Large RDMA Send• Small & Large RDMA Verify	

13.6: TI RDMA Stress

Test Result	Pass
Discussion:	
All switches were seen to properly handle a large load as indicated by the successful completion of control communications between two RNICs while other RNICs in the fabric were used to generate traffic in order to put a high load on the switch.	

13.7: TI Rsockets

Test Result	Pass
Discussion:	
Utilizing the rstream utility, all RCAs were able to successfully perform General, Asynchronous, Blocking, and Non-blocking procedures.	

13.8: TI MPI – Open MPI

Test Result	Pass
Discussion:	
Intel MPI Benchmarks were performed between all RCAs and were observed to exhibit the successful behavior.	