



OpenFabrics Alliance

Interoperability Logo Group (OFILG)

May 2014 Logo Event Report

UNH-IOL – 121 Technology Drive, Suite 2 – Durham, NH 03824 – +1-603-862-0090
OpenFabrics Interoperability Logo Group (OFILG) – ofalab@iol.unh.edu

Pradeep Satyanarayana
 IBM
 1385 NW Amberglen Parkway
 Hillsboro, OR 97006

Date: July 31, 2014
 Report Revision: 1.2
 OFED Version: 3.12
 OS Version: RHEL 6.4

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

IBM Power 7 R2 Server with RoCE Channel Adapter

The test suite referenced in this report is available at the UNH-IOL website. Release 1.50 (2014-May-6) was used.

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

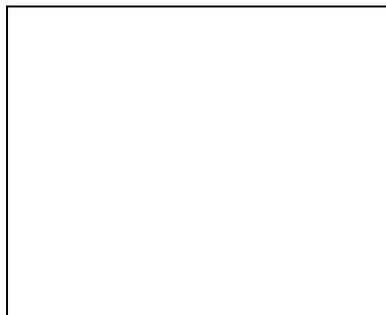
The following table highlights the Mandatory test results required for the OpenFabrics Interoperability Logo for the Server System using RoCE Channel Adapter and Running OFA Software device class per the Test Plan referenced above and the current OpenFabrics Interoperability Logo Program (OFILP).

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

Summary of all results follows on the second page of this report.
 For specific details regarding issues, please see the corresponding test result.

Testing Completed June 11, 2014

David Wyman
dwyman@iol.unh.edu



Review Completed July 31, 2014

Edward L. Mossman
emossm@iol.unh.edu

Result Summary

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

Test Procedures	IWG Test Status	Result/Notes
12.2: RoCE Link Initialization	Mandatory	PASS
12.4: RoCE IPoCE	Beta	PASS with Comments
13.2: TI NFS over RDMA	Beta	Not Applicable
13.4: TI uDAPL	Mandatory	PASS
13.5: TI RDMA Basic Interop	Mandatory	PASS
13.6: TI RDMA Stress	Beta	Not Tested
13.7: TI Rsockets	Mandatory	PASS
13.8: TI MPI – OpenMPI	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: 41 1E 00 9F 79 4D 02 EF E6 95 65 57 A4 71 4F 9F

SHA-1 Fingerprint: 44 51 9E 22 66 59 1A D3 A1 F9 0B EE BD 01 90 80 BE 61 A4 A8

Report Revision History

- v1.0 Initial Release
- v1.1 Modified IPoCE result to Pass with Comments
- v1.2 Modified device type from RoCE Adapter to Server System using RoCE Channel Adapter Running OFA Software

Configuration Files

Description	Attachment
RHEL 6.4 Configuration File	
OFED 3.12 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 currently 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.

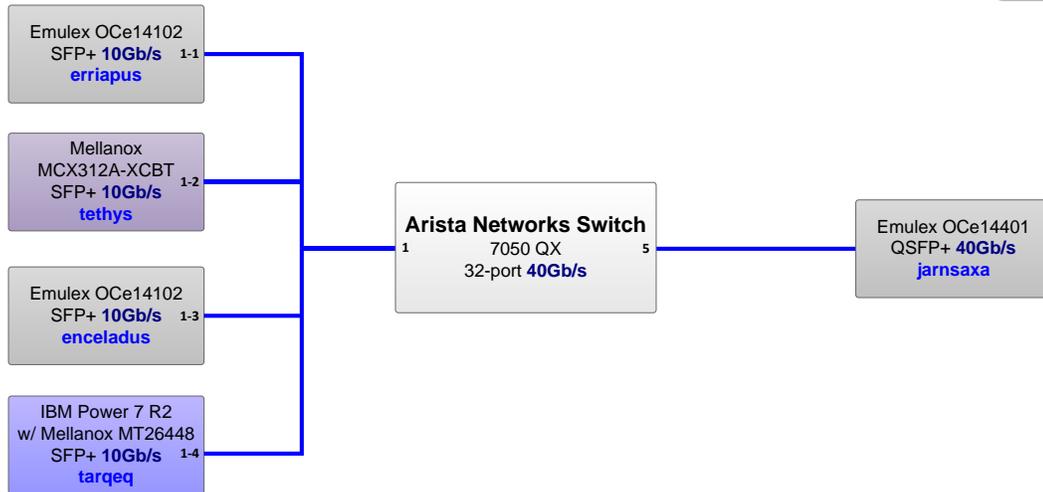
DUT and Test Setup Information

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

May 2014
 RoCE Topology

Ethernet Addressing
 <hostname>.ofa

RoCE Addressing
 <hostname>-ce.ofa



DUT Details			
Manufacturer:	IBM		
Machine:	IBM Power 7 R2		
Board:	OF30000010	Firmware Revision:	2.9.1200
Model:	Mellanox MT26448	Hardware Revision:	B0
Speed:	10Gb/s	Host:	tarqeq
Firmware MD5sum:			
Additional Comments / Notes:			

Mandatory Tests – RoCE Server System Test Results:

12.2: RoCE Link Initialization

Test Result	PASS
Result Discussion:	
All RoCE Channel Adapters were observed to link at the expected speed.	

Link Partner	IBM Power 7 R2
RCA: IBM Power 7 R2 - tarqeq	N/A
RCA: Emulex OCE14401 in host jarnsaxa	PASS
RCA: Emulex OCE14102 in host erriapus	PASS
RCA: Mellanox MCX312A in host methone	PASS

13.4: TI uDAPL

Test Result	PASS
Discussion:	
All of the Group 1: Point-to-Point and Group 2: Switched Topology tests were performed using OFED 3.12.	
Group 3: Switched Topology with Multiple Switches tests were not performed using OFED 3.12 due to time and equipment constraints.	

13.5: TI RDMA Basic Interoperability

Test Result	PASS
Discussion:	
RDMA read and write operations were observed to perform successfully between all RCAs in the cluster.	

13.6: TI RDMA Stress

Test Result	Not Tested
Discussion:	
This test requires a multi switch topology; it was therefore not performed due to equipment constraints.	

13.7: TI RSocket

Test Result	PASS
Discussion:	
All RCAs were able to successfully perform all Asynchronous, Blocking, and Non-blocking procedures.	

13.8: TI MPI – Open MPI

Test Result	PASS
Discussion:	
<p>The RCA in the IBM PPC64 system was excluded from the larger OpenMPI ring due to OpenMPI's limitations in regards to heterogeneous rings, and was instead tested against itself by being put in a "loopback" mode where it was able to complete all benchmarks successfully. This was accomplished by running the following command:</p> <pre>`/usr/local/bin/mpirun --mcabtolopenib,self --bind-to-core 16 --host tarqeq.ofa /usr/local/bin/IMB-MPI1`</pre>	

Beta Tests – RoCE Server System Test Results:

12.4: RoCE IPoCE

Test Section	Part A	Part B	Part C
Test Result	PASS	Not Available	PASS
Result Discussion:			
Part A: All devices were observed to successfully ping each Link Partner in the topology. Part B: Not available, as this test requires a second Ethernet switch which was not available in the topology. Part C: Every node was observed to correctly SFTP the test file to every other node in the topology.			

13.2: TI NFS over RDMA

Test Result	Not Applicable
Result Discussion:	
Due to ongoing debug and testing efforts related to OFED bug 2449 , the DUT was observed to not interoperate with Link Partners from other vendors.	