



OpenFabrics Alliance

Interoperability Logo Group (OFILG)

February 2017 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

Pradeep Satyanarayana
IBM
1385 NW Amberglen Parkway
Hillsboro, OR 97006

Date: February 15, 2017
Report Revision: 1.0
OFED Version: 3.18-2
OS Version: RHEL 7.2

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

IBM Power8 R2 Server with RoCE Channel Adapter

The test suite referenced in this report is available at the UNH-IOL website. Release 2.05 (2017-06-16) 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
12.4: IPoCE	Mandatory	Pass
13.2: TI NFSoRDMA	Beta	Pass
13.4: TI uDAPL	Mandatory	Pass
13.5: TI RDMA Basic Interop	Mandatory	Pass
13.6: TI RDMA Stress	Mandatory	Not Tested
13.7: TI RSocket	Mandatory	Pass
13.8: TI MPI – OpenMPI	Mandatory	Pass

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

Testing Completed February 17, 2017

Stefan Oesterreich
soesterreich@iol.unh.edu

Adam LeBlanc
aleblanc@iol.unh.edu

Reviewed & Issued February 17, 2017

Bob Noseworthy
ren@iol.unh.edu

Result Summary

The Following table summarizes all results from the event pertinent to this RoCE device class (Server Systems using RoCE Channel Adapters and running OFA software).

Test Procedures	IWG Test Status	Result/Notes
12.2: RoCE Link Initialization	Mandatory	Pass
12.4: IPoCE	Mandatory	Pass
13.2: TI NFSoRDMA	Beta	Pass
13.4: TI uDAPL	Mandatory	Pass
13.5: TI RDMA Basic Interop	Mandatory	Pass
13.6: TI RDMA Stress	Mandatory	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: 7B 9B 0C 40 55 27 86 C0 F7 4A A3 45 DB F9 40 6E
SHA-1 Fingerprint: 03 59 97 71 28 ED 17 7F 1A 83 C5 D0 1D A8 2B 98 3E 2F 0F E7

Report Revision History

- v1.0 Initial Release

Configuration Files

Description	Attachment
RHEL 7.2 Configuration File	
OFED 3.18-2 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.

February 2017 Logo
RoCE Topology

Ethernet Addressing
<hostname>.ofa

RoCE Addressing
<hostname>-ce.ofa



DUT Details			
Manufacturer:	IBM		
Machine:	IBM Power 8R2		
Board:	OF30000010	Firmware Revision:	2.35.5100
Model:	Mellanox MT27500	Hardware Revision:	0
Speed:	40Gb/s	Host:	Hyrrokkin-2
Firmware MD5sum:	N/A		
Additional Comments / Notes:			

Mandatory Tests –RoCE Server System Test Results:

12.2: RoCELink 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
Discussion:	
All devices were configured to have a IP addresses over their CR interface.	

13.2: TI NFSoRDMA

Test Result	Pass
Discussion:	
All DUT's were observed to successfully utilize NFS over RDMA.	

13.4: TI uDAPL

Test Result	Pass
Discussion:	
With the use of the daplttest utility, devices exhibited simple send receive, variation, 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	Not Tested
Discussion:	
Not run due to limited devices in the cluster. Minimum requirement states at least 4 devices are necessary.	

13.7: TI RSocket

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 – OpenMPI

Test Result	Pass
Discussion:	
MPI was run on a loop back interface since there are no other big endian machines in the cluster and was able to run the Intel MPI Benchmark successfully	