



OpenFabrics Alliance

Interoperability Logo Group (OFILG)

Dec 2011 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

Amit Krig
Mellanox Technologies, Ltd.
Beit Mellanox
Yokneam, Israel 20692

Date: 16 Mar 2012
Report Revision: 1.2
OFED Version on Compute Nodes: 1.5.4
Operating System on Compute Nodes: SL 6.1

Enclosed are the results from OFA Logo testing performed on the following devices under test (DUTs):
Mellanox SX5036

FDR devices are not yet supported by the logo program, but were tested at the December 2011 event and are listed in the Beta section for informative purposes. This is because the IBTA FDR Specification was not finalized in time for this event and none of the applicable test plan procedures had ever been run using FDR products.

The test suite referenced in this report is available at the IOL website. Release 1.4 (2011-Oct-25) was used.
http://www.iol.unh.edu/services/testing/ofa/testsuites/OFA-IWG_Interoperability_Test_Plan-v1.40.pdf

The logo document referenced in this report is available at the IOL website. Release 1.14 (2011-Mar-01) was used.
http://www.iol.unh.edu/services/testing/ofa/logoprogram/OFA-UNH-IOL_Logo_Program-v1.14.pdf

The Following Table highlights the Mandatory test results required for the OpenFabrics Interoperability Logo for the DUT per the Test Plan referenced above and the current OpenFabrics Interoperability Logo Program (OFILP).

Additional beta testing was performed using the DUT than is reflected in this report. A separate report will outline those results.

Test Procedures	IWG Test Status	Result/Notes
10.1: Link Initialization	Mandatory	FAIL
10.2: IB Fabric Initialization	Mandatory	PASS
10.3: IPoIB Connected Mode	Mandatory	PASS
10.4: IPoIB Datagram Mode	Mandatory	PASS
10.5: SM Failover and Handover	Mandatory	PASS
10.6: SRP	Mandatory	PASS
12.1: TI iSER	Mandatory	Not Available
12.2: TI NFS over RDMA	Mandatory	PASS
12.3: TI RDS	Mandatory	PASS
12.4: TI SDP	Mandatory	PASS
12.5: TI uDAPL	Mandatory	PASS
12.6: TI RDMA Basic Interop	Mandatory	PASS
12.8: TI RDMA Stress	Mandatory	PASS
12.11: TI MPI – Open	Mandatory	PASS
12.12: TI MPI – OSU	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 05 January 2012

Nickolas Wood
ndv2@iol.unh.edu



Review Completed 16 March 2012

Bob Noseworthy
ren@iol.unh.edu

Result Summary

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

Test Procedures	IWG Test Status	Result/Notes
10.1: Link Initialization	Mandatory	FAIL
10.2: IB Fabric Initialization	Mandatory	PASS
10.3: IPoIB Connected Mode	Mandatory	PASS
10.4: IPoIB Datagram Mode	Mandatory	PASS
10.5: SM Failover and Handover	Mandatory	PASS
10.6: SRP	Mandatory	PASS
10.7: Ethernet Gateway	Beta	Not Tested
10.8: FibreChannel Gateway	Beta	Not Tested
12.1: TI iSER	Mandatory	Not Available
12.2: TI NFS over RDMA	Mandatory	PASS
12.3: TI RDS	Mandatory	PASS
12.4: TI SDP	Mandatory	PASS
12.5: TI uDAPL	Mandatory	PASS
12.6: TI RDMA Basic Interoperability	Mandatory	PASS
12.8: TI RDMA Stress	Mandatory	PASS
12.10: TI MPI – Intel	Beta	Not Tested
12.11: TI MPI – Open	Mandatory	PASS
12.12: TI MPI – OSU	Mandatory	PASS

Digital Signature Information

This document was signed 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/certificates_and_fingerprints.php



If the document status still indicated "Validity of author NOT confirmed", then please contact the UNH-IOL to confirm the document's authenticity. To further validate the certificate integrity, Adobe 9.0 should report the following fingerprint information:

MD5 Fingerprint: B4 7E 04 FE E8 37 D4 D2 1A EA 93 7E 00 36 11 F3
SHA-1 Fingerprint: 50 E2 CB 10 21 32 33 56 4A FC 10 4F AD 24 6D B3 05 22 7C C0

Report Revision History

- v1.0 Initial working copy
- v1.1 Revised working copy
- v1.2 Post arbitration resolution update

Configuration Files

Description	Attachment
Scientific Linux 6.1 Configuration File	
OFED 1.5.4 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.
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 IB fabric configuration utilized for any tests requiring a multi-switch configuration is shown below.



DUT #1 Details			
Manufacturer:	Mellanox	Firmware Revision:	3.1.1080
Model:	SX6036	Hardware Revision:	B1
Speed:	FDR	Located in Host:	NA
Firmware MD5sum:	3bcb05600a667ca6c060a81191e5f847		
Additional Comments / Notes:			

Mandatory Tests – IB Device Test Results:

10.1: Link Initialization

Results	
Part #1:	FAIL
Discussion:	
It was noted during testing that certain port combinations between the Mellanox SX6036 switch and the QLogic 12200 switch linked at 4X DDR instead of the expected 4x QDR. This was more readily apparent when using passive copper cables five (5) meters in length when connected between ports on with long traces (generally edge ports).	

Link Partner		SX6036	SX6025
QLogic 12200 (Switch) – QDR		FAIL	PASS
Mellanox SX6025 (Switch) – FDR		PASS	NA
Mellanox SX6036 (Switch) – FDR		NA	PASS
Mellanox IS-5030 (Switch) – QDR		PASS	PASS
Mellanox BX5020 (Gateway) – QDR		PASS	PASS
DataDirect Networks SFA10000 (SRP Target) – QDR		PASS	PASS
DataDirect Networks S2A9900 (SRP Target) – DDR		PASS	PASS
NetApp Pikes Peak (SRP Target) – QDR		PASS	PASS
NetApp XBB2 (SRP Target) – DDR		PASS	PASS
Host: Themis	HCA: MHQH29C-XTR (QDR)	PASS	PASS
Host: Pan	HCA: MHQH19B-XTR (QDR)	PASS	PASS
Host: Hati	HCA: MCX353A-FCBT (FDR)	PASS	PASS
Host: Titan	HCA: MCX354A-FCBT (FDR)	PASS	PASS

10.2: Fabric Initialization

Subnet Manager				
OpenSM	IS-5030 SM	SX-6036 SM	QL12200 SM	WinOF SM
PASS	PASS	PASS	PASS	PASS
Result Discussion:				
All subnet managers used while testing with OFED 1.5.4 were able to correctly configure the selected topology.				

10.3: IPoIB Connected Mode

Subnet Manager					
Part	OpenSM	IS-5030 SM	SX-6036 SM	QL12200 SM	WinOF SM
A – ping	PASS	PASS	PASS	PASS	PASS
B – SFTP	PASS	PASS	PASS	PASS	PASS
C – SCP	PASS	PASS	PASS	PASS	PASS
Result Discussion:					
IPoIB ping, SFTP, and SCP transactions completed successfully between all HCAs; each HCA acted as both a client and a server for all tests.					

10.4: IPoIB Datagram Mode

Subnet Manager					
Part	OpenSM	IS-5030 SM	SX-6036 SM	QL12200 SM	WinOF SM
A – ping	PASS	PASS	PASS	PASS	PASS
B – SFTP	PASS	PASS	PASS	PASS	PASS
C – SCP	PASS	PASS	PASS	PASS	PASS
Result Discussion:					
IPoIB ping, SFTP, and SCP transactions completed successfully between all HCAs; each HCA acted as both a client and a server for all tests.					

10.5: SM Failover and Handover

SM Pairings		Result
OpenSM OFED 1.5.4	OpenSM OFED 1.5.4	PASS
Result Discussion:		
OpenSM was able to properly handle SM priority and state rules.		

10.6: SRP

Subnet Manager				
OpenSM	IS-5030 SM	SX-6036 SM	QL12200 SM	WinOF SM
PASS	PASS	PASS	PASS	PASS
Result Discussion:				
SRP communications between all HCAs and all SRP targets succeeded while the above mentioned SMs were in control of the fabric.				

12.1 TI iSER

Subnet Manager				
OpenSM	IS-5030 SM	SX-6036 SM	QL12200 SM	WinOF SM
Not Tested	Not Tested	Not Tested	Not Tested	Not Tested
Result Discussion:				
This test was not performed as there are no devices that support the iSER test procedure present in event topology.				

12.2: TI NFS over RDMA

Subnet Manager				
OpenSM	IS-5030 SM	SX-6036 SM	QL12200 SM	WinOF SM
PASS	PASS	PASS	PASS	PASS
Result Discussion:				
Connectathon was used to test NFS over RDMA; each HCA acted as both a client and a server.				

12.3: TI RDS

Subnet Manager					
Part	OpenSM	IS-5030 SM	SX-6036 SM	QL12200 SM	WinOF SM
A – ping	PASS	PASS	PASS	PASS	PASS
B – stress	PASS	PASS	PASS	PASS	PASS
Result Discussion:					
The reliable datagram socket protocol was tested between all HCAs; all communications completed successfully.					

12.4: TI SDP

Subnet Manager					
Part	OpenSM	IS-5030 SM	SX-6036 SM	QL12200 SM	WinOF SM
A – netperf	PASS	PASS	PASS	PASS	PASS
B – SFTP	PASS	PASS	PASS	PASS	PASS
C – SCP	PASS	PASS	PASS	PASS	PASS
Result Discussion:					
All communications using the SDP protocol completed successfully; each HCA acted as both a client and a server for all tests.					

12.5: TI uDAPL

Subnet Manager					
OpenSM	IS-5030 SM	SX-6036 SM	QL12200 SM	WinOF SM	
PASS	PASS	PASS	PASS	PASS	
Result Discussion:					
All communications using DAPL were seen to complete successfully as described in the referenced Test plan; each HCA acted as both a client and a server for all tests.					

12.6: TI RDMA Basic Interoperability

Subnet Manager					
OpenSM	IS-5030 SM	SX-6036 SM	QL12200 SM	WinOF SM	
PASS	PASS	PASS	PASS	PASS	
Result 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.					

12.8: TI RDMA Stress

Subnet Manager					
OpenSM	IS-5030 SM	SX-6036 SM	QL12200 SM	WinOF SM	
PASS	PASS	PASS	PASS	PASS	
Result Discussion:					
All IB switches were seen to properly handle a large load as indicated by the successfully completion of control communications between two HCAs while all other HCAs in the fabric were used to generate traffic in order to put a high load on the switch. Each HCA acted as both a client and a server for the control connection.					

12.11: TI MPI – Open

Subnet Manager					
Part	OpenSM	IS-5030 SM	SX-6036 SM	QL12200 SM	WinOF SM
A	PASS	PASS	PASS	PASS	PASS
B	PASS	PASS	PASS	PASS	PASS
Result Discussion:					
Complete heterogeneity; 1 MPI process per OFED 1.5.4 deployed system as described in the cluster topology (red and purple system icons), IB device vendor agnostic.					

12.12: TI MPI – OSU

Subnet Manager					
Part	OpenSM	IS-5030 SM	SX-6036 SM	QL12200 SM	WinOF SM
A	PASS	PASS	PASS	PASS	PASS
B	PASS	PASS	PASS	PASS	PASS
Result Discussion:					
Complete heterogeneity; 1 MPI process per OFED 1.5.4 deployed system as described in the cluster topology (red and purple system icons), IB device vendor agnostic.					

Beta Tests – IB Device Test Results:

10.7: IB Ethernet Gateway

Subnet Manager				
OpenSM	IS-5030 SM	SX-6036 SM	QL12200 SM	WinOF SM
Not Tested	Not Tested	Not Tested	Not Tested	Not Tested
Result Discussion:				
This test was not performed as there are no devices that support the Ethernet Gateway test procedure present in event topology.				

10.8 IB FibreChannel Gateway

Subnet Manager				
OpenSM	IS-5030 SM	SX-6036 SM	QL12200 SM	WinOF SM
Not Tested	Not Tested	Not Tested	Not Tested	Not Tested
Result Discussion:				
This test was not performed as there are no devices that support the FibreChannel Gateway test procedure present in event topology.				

12.10: MPI – Intel

Subnet Manager				
OpenSM	IS-5030 SM	SX-6036 SM	QL12200 SM	WinOF SM
Not Tested	Not Tested	Not Tested	Not Tested	Not Tested
Result Discussion:				
This test was not performed as the binaries for Intel MPI are not present on the compute nodes present in event topology.				