



OpenFabrics Alliance

Interoperability Logo Group (OFILG)

May 2012 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

Idan Kligvasser
Mellanox Technologies
Hermon Building 4th Floor
P.O. Box 586, Yokenam 20692
Israel

Date: 9 August 2012
Report Revision: 2.0
OFED Version on Compute Nodes: 1.5.4.1
Operating System on Compute Nodes: SL 6.2

Enclosed are the results from OFA Logo testing performed on the following devices under test (DUTs):
Mellanox MCX353A-FCBT Mellanox MCX354A-FCBT

The test suite referenced in this report is available at the IOL website. Release 1.42 (2012-Apr-03) was used.

[www.iol.unh.edu/services/testing/ofa/testsuites/OFA-IWG Interoperability Test Plan-v1.42.pdf](http://www.iol.unh.edu/services/testing/ofa/testsuites/OFA-IWG_Interoperability_Test_Plan-v1.42.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 testing was performed using the DUT then is reflected in this report. A separate report will outline those results.

Test Procedures	IWG Test Status	Result/Notes
10.1: Link Initialization	Mandatory	PASS
10.2: 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	Not Tested
12.3: TI RDS	Mandatory	PASS
12.4: TI SDP	Mandatory	PASS
12.5: TI uDAPL	Mandatory	PASS with Comments
12.6: TI RDMA Basic Interoperability	Mandatory	PASS
12.8: TI RDMA Stress	Mandatory	PASS
12.11: TI MPI – Open	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 09 August 2012

Edward L. Mossman
emossm@iol.unh.edu



Review Completed 10 August 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	PASS
10.2: 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	Not Tested
12.3: TI RDS	Mandatory	PASS
12.4: TI SDP	Mandatory	PASS
12.5: TI uDAPL	Mandatory	PASS with Comments
12.6: TI RDMA Basic Interoperability	Mandatory	PASS
12.8: TI RDMA Stress	Mandatory	PASS
12.11: TI MPI – Open	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
- v2.0 Resolved uDAPL failure through arbitration.

Configuration Files

Description	Attachment
Scientific Linux 6.2 Configuration File	
OFED 1.5.4.1 Configuration File	

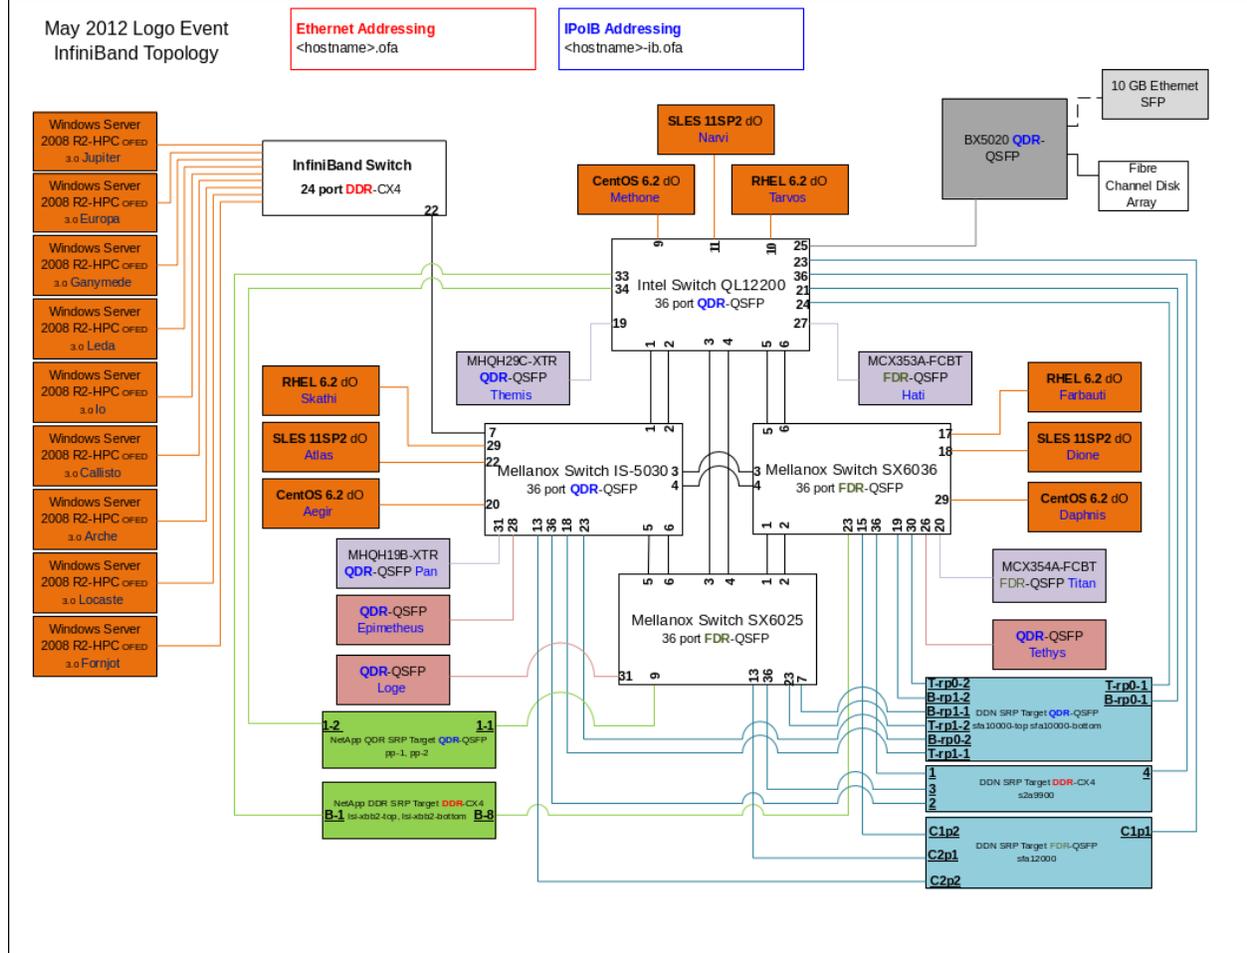
Result Key

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.
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:	2.10.4530
Model:	MCX353A-FCBT	Hardware Revision:	X2
Speed:	FDR	Located in Host:	Hati
Firmware MD5sum:	da6c7ae3c5da51301e93cd152492d611		
Additional Comments / Notes:			

DUT #2 Details			
Manufacturer:	Mellanox	Firmware Revision:	2.10.4530
Model:	MCX354A-FCBT	Hardware Revision:	X2
Speed:	FDR	Located in Host:	Titan
Firmware MD5sum:	0a2205eacd95354845ad527bd4b34686		
Additional Comments / Notes:			

Mandatory Tests – IB Device Test Results:

10.1: Link Initialization

Results	
Part #1:	PASS
Discussion:	
All links established with the DUT were of the proper link speed and width.	

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

10.2: Fabric Initialization

Subnet Manager				
OpenSM	IS-5030 SM	SX-6036 SM	12200 SM	WinOF SM
PASS	PASS	PASS	PASS	PASS
Result Discussion:				
All subnet managers used while testing with OFED 1.5.4.1 were able to correctly configure the selected topology.				

10.3: IPoIB Connected Mode

Subnet Manager					
Part	OpenSM	IS-5030 SM	SX-6036 SM	12200 SM	WinOF SM
A	PASS	PASS	PASS	PASS	PASS
B	PASS	PASS	PASS	PASS	PASS
C	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	12200 SM	WinOF SM
A	PASS	PASS	PASS	PASS	PASS
B	PASS	PASS	PASS	PASS	PASS
C	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 1.5.4.1	OpenSM OFED 1.5.4.1	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	12200 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	12200 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 the event topology.				

12.2: TI NFS over RDMA

Subnet Manager				
OpenSM	IS-5030 SM	SX-6036 SM	12200 SM	WinOF SM
Not Tested	Not Tested	Not Tested	Not Tested	Not Tested
Result Discussion:				
NFS over RDMA is not supported in the version of the Linux kernel used during this event (2.6.32); therefore this test could not be performed.				

12.3: TI RDS

Subnet Manager					
Part	OpenSM	IS-5030 SM	SX-6036 SM	12200 SM	WinOF SM
A	PASS	PASS	PASS	PASS	PASS
B	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	12200 SM	WinOF SM
A	PASS	PASS	PASS	PASS	PASS
B	PASS	PASS	PASS	PASS	PASS
C	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	12200 SM	WinOF SM
PASS with Comments	PASS with Comments	PASS with Comments	PASS with Comments	PASS with Comments
Result Discussion:				
<p>The DUTs were observed to timeout during tests 2.2 and 3.1 when connected via a cable rated for FDR speeds. When using a cable rated for QDR speeds, no timeouts were observed. This suggests an incompatibility with the daplttest utility and devices operating at FDR speeds.</p> <p>Upon upgrading to OFED 3.5-rc1, no timeouts were observed. OFED 3.5-rc1 contains an updated version of the daplttest utility, leading to the conclusion that the version of daplttest contained in OFED 1.5.4.1 incorrectly supported devices operating at FDR speeds.</p>				

12.6: TI RDMA Basic Interoperability

Subnet Manager				
OpenSM	IS-5030 SM	SX-6036 SM	12200 SM	WinOF SM
PASS	PASS	PASS	PASS	PASS
Result Discussion:				
<p>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.</p>				

12.8: TI RDMA Stress

Subnet Manager				
OpenSM	IS-5030 SM	SX-6036 SM	12200 SM	WinOF SM
PASS	PASS	PASS	PASS	PASS
Result Discussion:				
<p>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 acted as noise on the fabric. Each HCA acted as both a client and a server for the control connection.</p>				

12.11: TI MPI – Open

	Subnet Manager				
Part	OpenSM	IS-5030 SM	SX-6036 SM	12200 SM	WinOF SM
A	PASS	PASS	PASS	PASS	PASS
B	PASS	PASS	PASS	PASS	PASS
Result Discussion:					
<p>Complete heterogeneity; one process per system as described in the cluster topology.</p>					