



OpenFabrics Alliance

Interoperability Logo Group (OFILG)

July 2015 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

Guy Ergas
Mellanox Technologies
Beit Mellanox, 2nd Floor
P.O. Box 586, Yokenam 20692
Israel

Date: 31 August 2015
Report Revision: 1.0
OFED Version on Compute Nodes: 3.18
Operating System on Compute Nodes: Scientific Linux 7.0

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

Mellanox SX6036
Mellanox IS-5030

Mellanox SX6025
Mellanox SX6036G

The test suite referenced in this report is available at the UNH-IOL website. Release 2.04 (2015-July-23) was used.

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

The following table highlights the Mandatory tests required for the OpenFabrics Interoperability Logo for the InfiniBand Switch device class per the Test Plan & the current OpenFabrics Interoperability Logo Program (OFILP).

Test Procedures	IWG Test Status	Result/Notes
11.1: Link Initialization	Mandatory	PASS
11.2: Fabric Initialization	Mandatory	PASS
11.3: IPoIB Connected Mode	Mandatory	PASS
11.4: IPoIB Datagram Mode	Mandatory	PASS with Comments
11.5: SM Failover and Handover	Mandatory	PASS
11.6: SRP	Mandatory	PASS
13.2: TI NFS over RDMA	Mandatory	PASS
13.4: TI uDAPL	Mandatory	PASS
13.5: TI RDMA Basic Interoperability	Mandatory	PASS
13.6: TI RDMA Stress	Mandatory	PASS
13.7: TI Rsockets	Mandatory	PASS
13.8: 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 August 31, 2015

Stefan Oesterreich
soesterreich@iol.unh.edu



Reviewed & Issued October 2, 2015

Bob Noseworthy
ren@iol.unh.edu

Result Summary

The Following table summarizes all results from the event pertinent to this IB device class (InfiniBand Switch).

Test Procedures	IWG Test Status	Result/Notes
11.1: Link Initialization	Mandatory	PASS
11.2: Fabric Initialization	Mandatory	PASS
11.3: IPoIB Connected Mode	Mandatory	PASS
11.4: IPoIB Datagram Mode	Mandatory	PASS with Comments
11.5: SM Failover and Handover	Mandatory	PASS
11.6: SRP	Mandatory	PASS
11.7: IB Ethernet Gateway	Beta	Not Tested
11.8: IB FibreChannel Gateway	Beta	Not Tested
13.2: TI NFS over RDMA	Mandatory	PASS
13.4: TI uDAPL	Mandatory	PASS
13.5: TI RDMA Basic Interoperability	Mandatory	PASS
13.6: TI RDMA Stress	Mandatory	PASS
13.7: TI Rsockets	Mandatory	PASS
13.8: 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/>



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: FF 91 7B BD 2E 1A 0E 24 16 A8 23 28 13 69 D0 72
SHA-1 Fingerprint: 0C 88 5A 63 08 51 9B E0 D1 96 59 62 5E B3 52 01 58 C9 AF 27

Report Revision History

- v1.0 Initial working copy

Configuration Files

Description	Attachment
Scientific Linux 7.0 Configuration File	
OFED 3.18 Configuration File	

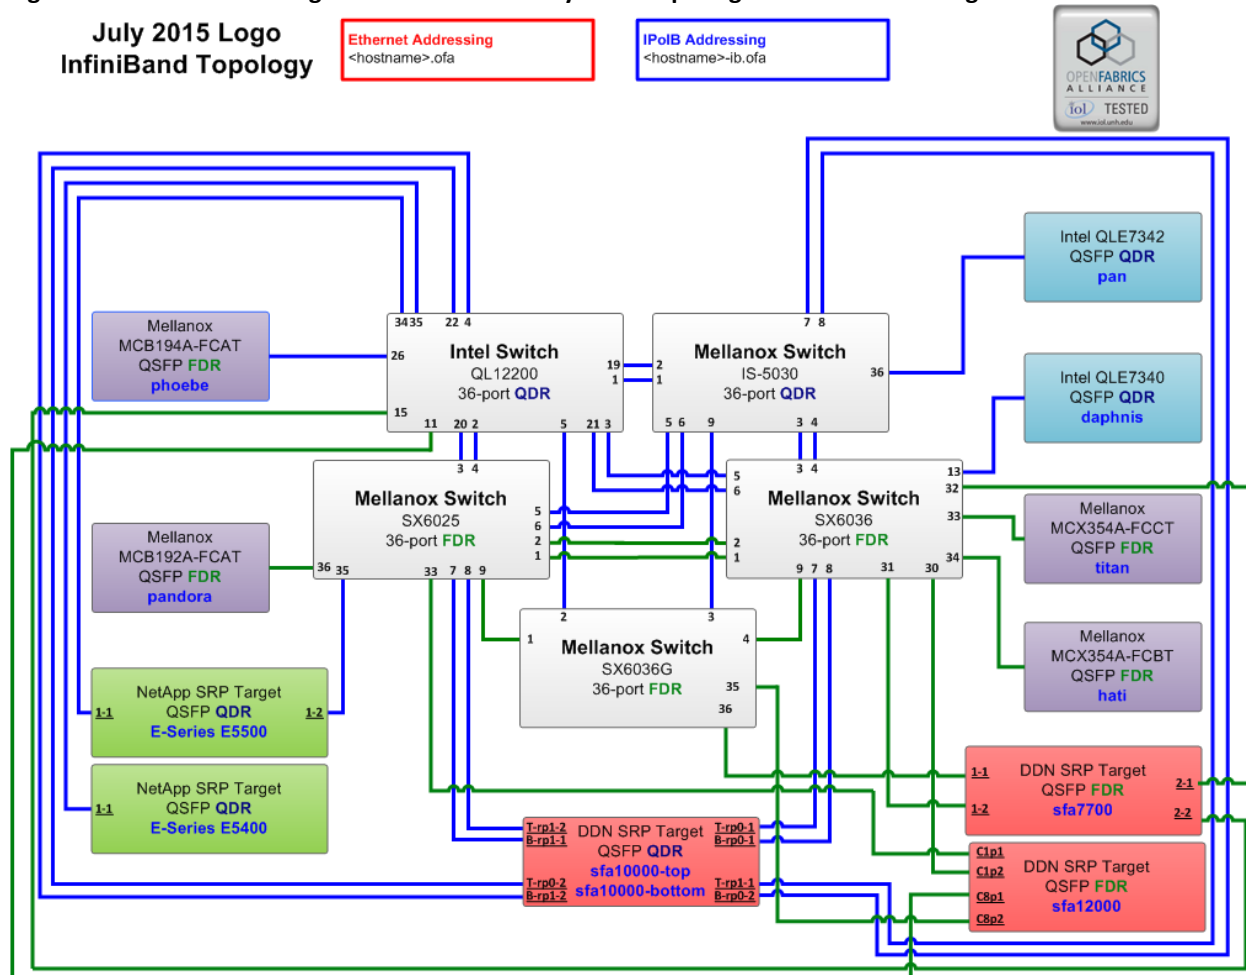
Result Key

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 IB fabric configuration utilized for any tests requiring a multi-switch configuration is shown below.



DUT #1 Details			
Manufacturer:	Mellanox	Firmware Revision:	9.3.1750
Model:	SX6036	Hardware Revision:	X2
Speed:	FDR	Located in Host:	N/A
Firmware MD5sum:	N/A		
Additional Comments / Notes:			
Version summary: PPC_M460EX 3.4.1110 2015-03-08 15:10:09 ppc			

DUT #2 Details			
Manufacturer:	Mellanox	Firmware Revision:	9.2.8000
Model:	SX6025	Hardware Revision:	X2
Speed:	FDR	Located in Host:	N/A
Firmware MD5sum:	N/A		
Additional Comments / Notes:			

DUT #3 Details			
Manufacturer:	Mellanox	Firmware Revision:	7.4.2360
Model:	IS-5030	Hardware Revision:	X2
Speed:	QDR	Located in Host:	N/A
Firmware MD5sum:	N/A		
Additional Comments / Notes:			
Version summary: EFM_PPC_M405EX EFM_1.1.3000 2013-07-08 14:29:44 ppc			

DUT #4 Details			
Manufacturer:	Mellanox	Firmware Revision:	9.3.1750
Model:	SX6036g	Hardware Revision:	
Speed:	FDR	Located in Host:	N/A
Firmware MD5sum:	N/A		
Additional Comments / Notes:			
Version summary: PPC_M460EX 3.4.1110 2015-03-08 15:10:09 ppc			

Mandatory Tests – IB Device Test Results:

11.1: Link Initialization

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

Link Partner	SX6025	SX6036	IS-5030	SX6036g	
QLogic 12200 (Switch) – QDR	PASS	PASS	PASS	PASS	
Mellanox SX6025 (Switch) – FDR	NA	PASS	PASS	PASS	
Mellanox SX6036 (Switch) – FDR	PASS	NA	PASS	PASS	
Mellanox IS-5030 (Switch) – QDR	PASS	PASS	NA	PASS	
Mellanox SX6036g (Switch) – FDR	PASS	PASS	PASS	NA	
DataDirect Networks SFA10000 (SRP Target) – QDR	PASS	PASS	PASS	PASS	
DataDirect Networks SFA12000 (SRP Target) – FDR	PASS	PASS	PASS	PASS	
DataDirect Networks SFA7700 (SRP Target) – FDR	PASS	PASS	PASS	PASS	
NetApp Soyuz (SRP Target) – QDR	PASS	PASS	PASS	PASS	
NetApp Pikes Peak (SRP Target) – QDR	PASS	PASS	PASS	PASS	
Host: hati	HCA: MCX354A-FCBT (FDR)	PASS	PASS	PASS	PASS
Host: titan	HCA: MCX354A-FCCT (FDR)	PASS	PASS	PASS	PASS
Host: phoebe	HCA: MCB194A-FCAT (FDR)	PASS	PASS	PASS	PASS
Host: pandora	HCA: MCB192A-FCAT (FDR)	PASS	PASS	PASS	PASS
Host: pan	HCA: QLE7342 (QDR)	PASS	PASS	PASS	PASS
Host: daphnis	HCA: QLE7340 (QDR)	PASS	PASS	PASS	PASS

11.2: Fabric Initialization

Subnet Manager	Result
OpenSM	PASS
Result Discussion:	
All subnet managers used while testing with OFED 3.12-1 were able to correctly configure the selected topology.	

11.3: IPoIB Connected Mode

Subnet Manager	Part A	Part B	Part C
OpenSM	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			

11.4: IPoIB Datagram Mode

Subnet Manager	Part A	Part B	Part C
OpenSM	PASS with Comments	PASS	PASS
Result Discussion:			
<p>The DUT has been noted not to receive the echo reply of the first ping sequence when the entry for the device is cleared in the ARP table and the payload size is 32768 or higher. To reproduce the issue, follow the testplan IPoIB section 11.4.1 Step 7 on page 61. The problem will occur upon execution of Step 7, just after the paired device’s InfiniBand interface has been removed from the ARP table. If the command run in Step 7, is run a subsequent time it is noted that all pings execute as expected and no problems are present. Thus, ping has been observed to work as expected and no problems are observed when the ARP table has been populated with an entry for the DUT.</p> <p>All permutations of this test on other devices in the test cluster were also noted to encounter the same issue. Therefore this issue is not currently attributed to this DUT and as such is noted here as a Pass with Comments. This may be attributed to an unresolved issue with OFED or Scientific Linux 7.0. These commands and sizes were observed to work with SL 6.5 and previous versions of OFED on previous devices.</p>			

11.5: SM Failover and Handover

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

11.6: SRP

Subnet Manager	Result
OpenSM	PASS
Result Discussion:	
SRP communications between all HCAs and all SRP targets succeeded while OpenSM was in control of the fabric.	

13.2: TI NFS over RDMA

Subnet Manager	Result
OpenSM	Refer to Comments

Result Discussion:

All DUT's were observed to successfully utilize NFS over RDMA.

13.4: TI uDAPL

Subnet Manager	Result
OpenSM	PASS
Result Discussion:	
All communications using uDAPL 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.	

13.5: TI RDMA Basic Interoperability

Subnet Manager	Result
OpenSM	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.	

13.6: TI RDMA Stress

Subnet Manager	Result
OpenSM	PASS
Result Discussion:	
All IB switches were seen to properly handle a large load as indicated by the successful 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.	

13.7: TI Rsockets

Subnet Manager	Result
OpenSM	PASS with Comments
Result Discussion:	
DUT's were observed to pass all rsockets procedures.	

13.8: TI MPI – Open

Subnet Manager	Part A	Part B
OpenSM	PASS	PASS
Result Discussion:		
HCAs were capable of running the mpirun binary in accordance to the current test plan between all other hosts.		