



# OpenFabrics Alliance

## Interoperability Working Group (OFA-IWG)

### June 2009 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  
 Hermon Building 4th Floor  
 P.O. Box 586, Yokenam 20692  
 Israel

August 21, 2009  
 Report Rev1.4  
 OFED Version: 1.4.1

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

*Mellanox MTS-3600 (36 Port QDR InfiniBand Switch)*

The test suite referenced in this report is available at the OFA website, at test time release 1.28 (March 27, 2009 DRAFT) was used:

[http://openfabrics.org/downloads/OFA-IWG Interoperability Test Plan-v1.28.pdf](http://openfabrics.org/downloads/OFA-IWG%20Interoperability%20Test%20Plan-v1.28.pdf)


The following table highlights the Mandatory test results required for the OpenFabrics Interoperability Logo for the DUT per the testplan referenced above and the current OpenFabrics Interoperability Logo Program (OFILP).

Mandatory Test Procedures	IWG Test Status	Result/Notes
10.1: IB Link Initialization	<b>Mandatory</b>	<b>PASS</b>
10.2: IB Fabric Initialization	<b>Mandatory</b>	<b>PASS</b>
10.3: IB IPoIB Connected Mode	<b>Mandatory</b>	<b>PASS</b>
10.4: IB IPoIB Datagram Mode	<b>Mandatory</b>	<b>PASS with Comments</b>
10.6: IB SRP	<b>Mandatory</b>	<b>PASS</b>
12.1: TI iSER	<b>Mandatory</b>	<b>Not Available</b>
12.4: TI SDP	<b>Mandatory</b>	<b>PASS</b>

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


Summary of all results follows on the second page of this report.

Testing Completed 06/19/2009

  
 Nickolas Wood  
 ndv2@iol.unh.edu



Review Completed 08/12/2009

  
 Bob Noseworthy  
 ren@iol.unh.edu

## Table 1: Result Summary

The following table summarizes all results from the event pertinent to an IB device.

Note: **Failing a beta test does not necessarily indicate a failure of the device under test.** There are several reasons that a device might fail a beta test that are beyond the control of the device manufacturer. Such reasons include unresolved problems with the OFED software, with other vendors' devices, and with the test plan procedures.

Test Procedures	IWG Test Status	Result/Notes
10.1: IB Link Initialization	<b>Mandatory</b>	<b>PASS</b>
10.2: IB Fabric Initialization	<b>Mandatory</b>	<b>PASS</b>
10.3: IB IPoIB Connected Mode	<b>Mandatory</b>	<b>PASS</b>
10.4: IB IPoIB Datagram Mode	<b>Mandatory</b>	<b>PASS with Comments</b>
10.5: IB SM Failover and Handover	Beta	<b>PASS</b>
10.6: IB SRP	<b>Mandatory</b>	<b>PASS</b>
10.7: IB Ethernet Gateway	Beta	<b>Not Applicable</b>
10.8: IB FibreChannel Gateway	Beta	<b>Not Applicable</b>
12.1: TI iSER	<b>Mandatory</b>	<b>Not Available</b>
12.10: HP MPI - HP	Beta	<b>Not Applicable</b>
12.11: TI MPI - Intel	Beta	<b>Not Applicable</b>
12.12: TI MPI - Open	Beta	<b>Not Applicable</b>
12.13: TI MPI - OSU	Beta	<b>Not Applicable</b>
12.3: TI Reliable Datagram Service	Beta	<b>Not Applicable</b>
12.4: TI SDP	<b>Mandatory</b>	<b>PASS</b>
12.5: TI uDAPL	Beta	<b>Not Applicable</b>
12.6-7: TI Basic RDMA Interoperability	Beta	<b>Not Tested</b>
12.8-9: TI RDMA Operations.	Beta	<b>Not Tested</b>

### *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/certificates\\_and\\_fingerprints.php](http://www.iol.unh.edu/certifyDoc/certificates_and_fingerprints.php)

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 should report the following fingerprint information:

MD5 Fingerprint: F6E2 1B99 28AD 0D25 E77E ADE5 479A 1E05

SHA-1 Fingerprint: AD30 8B08 DD3B B2E3 9362 46E9 3427 BE47 1D49 890B

## Report Revision History

- v1.0 Initial working copy
- v1.1 Fixed Device names and topology
- v1.2 Fixed device names and topology. Correct port count (to 36) on DUT. Added note about beta tests in results summary.
- v1.3 Changed topology diagram.
- v1.4 Add firmware version

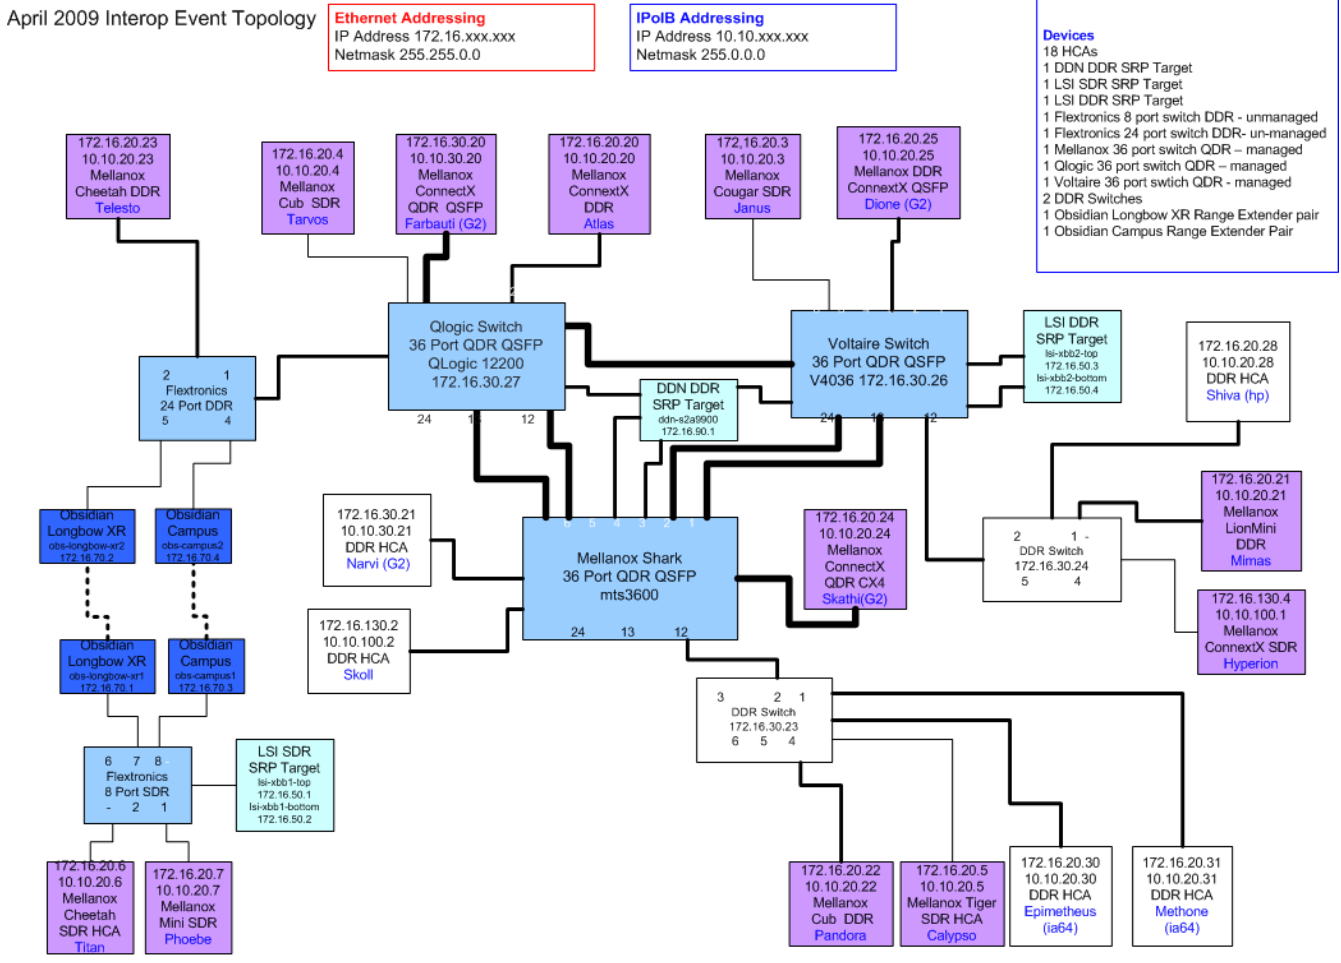
## Table 2: Result Key

The following table contains possible results and their meanings:

Result:	Description:
<b>PASS</b>	The Device Under Test (DUT) was observed to exhibit conformant behavior.
<b>PASS with Comments</b>	The DUT was observed to exhibit conformant behavior however an additional explanation of the situation is included, such as due to time limitations only a portion of the testing was performed.
<b>FAIL</b>	The DUT was observed to exhibit non-conformant behavior.
<b>Warning</b>	The DUT was observed to exhibit behavior that is not recommended.
<b>Informative</b>	Results are for informative purposes only and are not judged on a pass or fail basis.
<b>Refer to Comments</b>	From the observations, a valid pass or fail could not be determined. An additional explanation of the situation is included.
<b>Not Applicable</b>	The DUT does not support the technology required to perform this test.
<b>Not Available</b>	Due to testing station limitations or time limitations, the tests could not be performed.
<b>Borderline</b>	The observed values of the specified parameters are valid at one extreme and invalid at the other.
<b>Not Tested</b>	Not tested due to the time constraints of the test period.

### Table 3: 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 Rev	7.2.0
Model	MTS-3600	Hardware Rev	N/A
Speed	QDR 4x	IP Address in Fabric	N/A
Additional Comments/Notes			
Mellanox MTS-3600( 36 Port QDR InfiniBand Switch) with QSFP Ports			

## Mandatory Tests - IB Device Test Summary Results:

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 10.1: IB Link Initialize	Test #1	Phy link is established	<b>PASS</b>
<b>Discussion:</b>			
Physical link initialization was verified between this device and every other device in the fabric. Link status was observed visually via status lights on the device. Port width and link speed was verified via ibdiagnet.			
The MTS-3600 successfully links with all devices at the proper speed and port width.			

Link Partner Device	MTS-3600	
QLogic 12200 (Switch)	<b>PASS</b>	
Flextronics F-X430066 (Switch)	<b>PASS</b>	
Flextronics F-X430044 (Switch)	<b>PASS</b>	
Mellanox MTS3600 (Switch)	<b>N/A</b>	
Voltaire v4036 (Switch)	<b>PASS</b>	
Obsidian Longbow-XR (Range Extender)	<b>PASS</b>	
Obsidian Longbow-XR (Range Extender)	<b>PASS</b>	
Obsidian Longbow Campus (Range Extender)	<b>PASS</b>	
Obsidian Longbow Campus (Range Extender)	<b>PASS</b>	
LSI XBB1 (SRP Target)	<b>PASS</b>	
LSI XBB2 (SRP Target)	<b>PASS</b>	
DataDirect Networks (SRP Target)	<b>PASS</b>	
Host: Dione, G2 PCI e	HCA: Mellanox ConnectX DDR	<b>PASS</b>
Host: Hyperion	HCA: Mellanox LionCub DDR	<b>PASS</b>
Host: Skoll	HCA: DDR HCA	<b>PASS</b>
Host: Narvi, G2 PCI e	HCA: DDR HCA	<b>PASS</b>
Host: Telesto	HCA: Mellanox Cheetah DDR	<b>PASS</b>
Host: Pandora	HCA: Mellanox ConnectX SDR	<b>PASS</b>
Host: Titan	HCA: Mellanox Cheetah SDR	<b>PASS</b>
Host: Phoebe	HCA: Mellanox Lion Mini SDR	<b>PASS</b>
Host: Calypso	HCA: Mellanox Tiger SDR	<b>PASS</b>
Host: Skathi	HCA: Mellanox ConnectX QDR	<b>PASS</b>
Host: Janus	HCA: Mellanox Cougar SDR	<b>PASS</b>
Host: Tarvos	HCA: Mellanox Lion Cub SDR	<b>PASS</b>
Host: Farbauti, G2 PCI e	HCA: Mellanox ConnectX QDR	<b>PASS</b>
Host: Atlas	HCA: Mellanox ConnectX DDR	<b>PASS</b>
Host: Mimas	HCA: Mellanox Lion Mini DDR	<b>PASS</b>
Host: Shiva	HCA: Mellanox ConnectX DDR	<b>PASS</b>
Host: Methone, IA_64	HCA: Mellanox InfiniHost III DDR	<b>PASS</b>
Host: Epimetheus, IA_64	HCA: Mellanox InfiniHost III DDR	<b>PASS</b>

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 10.2: IB Fabric Initialization	Test #1:	Port is Active with all SMs	<b>PASS</b>

**Discussion:**

The fabric configuration shown in Figure 1 was used for this test. 'ibdiagnet -c 1000' showed no Port error counters increment. Only one SM is run at a time. All switches are power cycled between SM trials. All links are validated via use of 'ibdiagnet' and 'ibchecknet' was used to verify that there were no duplicate guides. Refer to the table below for SM details.

SMs tested include: *OFED OpenSM (SM Only)*, *QLogic 12200 (Managed Switch)*, *Mellanox MTS3600 (Managed Switch)* and *Voltaire v4036 (Managed Switch)*

<i>MTS-3600</i>	All ports Armed/Active	No Dup GUIDs	No Port errors
OFED OpenSM (SM only)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
QLogic 12200 (Managed Switch)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Mellanox MTS3600 (Managed Switch)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Voltaire v4036 (Managed Switch)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>

<i>MTS-3600</i>	All ports Armed/Active	No Dup GUIDs	No Port errors
OFED OpenSM (SM only)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
QLogic 12200 (Managed Switch)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Mellanox MTS3600 (Managed Switch)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Voltaire v4036 (Managed Switch)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 10.3: IPoIB Connected Mode	Test #1	Automated Test Script	<b>PASS</b>
	Test #2	Automated Test Script	<b>PASS</b>
	Test #3	Automated Test Script	<b>PASS</b>
<b>Discussion:</b>			
<p><b>Test #1:</b> An automated test script was used to send ICMP Echo Request packets with payloads of specific sizes between all hosts on the configured fabric. This procedure was repeated with each subnet manager independently managing the fabric.</p> <p><b>Test #2:</b> An HCA was disconnected from the fabric and reconnected in a different location; the ICMP Echo Reply packets ceased while the HCA was disconnected, and then resumed when it was reconnected. This procedure was repeated once with each subnet manager independently managing the fabric.</p> <p><b>Test #3:</b> An automated test script was used to transfer a 4MB file using the SFTP protocol between all hosts on the configured fabric. The file was transferred once in each direction between all hosts, and the contents of the file were verified after each transfer. This procedure was repeated with each subnet manager independently managing the fabric.</p> <p>SMs tested include: <i>OFED OpenSM (SM Only)</i>, <i>QLogic 12200 (Managed Switch)</i>, <i>Mellanox MTS3600 (Managed Switch)</i> and <i>Voltaire v4036 (Managed Switch)</i></p>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 10.4: IPoIB Datagram Mode	Test #1	Automated Test Script	<b>PASS with Comments</b>
	Test #2	Automated Test Script	<b>PASS</b>
	Test #3	Automated Test Script	<b>PASS</b>
<b>Discussion:</b>			
<p><b>Test #1:</b> An automated test script was used to send ICMP Echo Request packets with payloads of specific sizes between all hosts on the configured fabric. This procedure was repeated with each subnet manager independently managing the fabric.</p> <p>The last packet size tested; 65507 bytes, was seen to lose the very first packet sent. This can be fixed by setting the systems qlen value to 18 and as such has no bearing on this device.</p> <p><b>Test #2:</b> An HCA was disconnected from the fabric and reconnected in a different location; the ICMP Echo Reply packets ceased while the HCA was disconnected, and then resumed when it was reconnected. This procedure was repeated once with each subnet manager independently managing the fabric.</p> <p><b>Test #3:</b> An automated test script was used to transfer a 4MB file using the SFTP protocol between all hosts on the configured fabric. The file was transferred once in each direction between all hosts, and the contents of the file were verified after each transfer. This procedure was repeated with each subnet manager independently managing the fabric.</p> <p>SMs tested include: <i>OFED OpenSM (SM Only)</i>, <i>QLogic 12200 (Managed Switch)</i>, <i>Mellanox MTS3600 (Managed Switch)</i> and <i>Voltaire v4036 (Managed Switch)</i></p>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 10.6: IB SRP	Test #1:	Automated Test Script	<b>PASS</b>
<b>Discussion:</b>			
No SRP issues were observed with this switch.			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.1: TI iSER	Test #1-4	Not Available	<b>Not Available</b>
<b>Discussion:</b>			
This test was not performed due to the unavailability of an iSER target device in the fabric			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.4: TI SDP	Test #1: Netperf	Automated Test Script	<b>PASS</b>
	Test #2: SFTP	Automated Test Script	<b>PASS</b>
	Test #3: SCP	Automated Test Script	<b>PASS</b>
<b>Discussion:</b>			
No SDP issues were observed with this switch.			



## Beta Tests - IB Device Test Results:

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 10.5: IB SM Failover and Handover	Test #1	No problems detected	<b>PASS</b>
<b>Discussion: Test Results</b>			
SM failover/handover succeeded in a fabric utilizing these DUT's.			
OpenSM was the only SM tested.			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 10.7: IB Ethernet Gateway	Test #1-10	Not applicable to DUT	<b>Not Applicable</b>
<b>Discussion:</b>			
The OFA Logo Program does not require these tests to be performed on a device that does not support this functionality.			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 10.8: IB FibreChannel Gateway	Test #1-10	Not applicable to DUT	<b>Not Applicable</b>
<b>Discussion:</b>			
The OFA Logo Program does not require these tests to be performed on a device that does not support this functionality.			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.3: TI Reliable Datagram Service	Test #1-10	Not applicable to DUT	<b>Not Applicable</b>
<b>Discussion:</b>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.5: TI uDAPL	Test #1	Not applicable to DUT	<b>Not Applicable</b>
<b>Discussion:</b>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.6: TI Basic RDMA Interoperability – Using XanStorm	Test #1-10	Not Tested	<b>Not Tested</b>
<b>Discussion:</b>			
Due to time constraints this test was not performed.			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.7: TI Basic RDMA Interoperability – Using Command Line	Test #1-10	Not Tested	<b>Not Tested</b>
<b>Discussion:</b>			
Due to time constraints this test was not performed.			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.8: TI RDMA Operations. Using XANstorm	Test #1-10	Not Tested	<b>Not Tested</b>
<b>Discussion:</b>			
Due to time constraints this test was not performed.			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.9: TI RDMA Operations. Using Command Line	Test #1-10	Not Tested	<b>Not Tested</b>
<b>Discussion:</b>			
Due to time constraints this test was not performed.			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.10: MPI – Hewlett-Packard	Test #1-24	Not applicable to DUT	<b>Not Applicable</b>
<b>Discussion:</b>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.11: TI MPI – Intel	Test #1 PingPing and PingPong	Not applicable to DUT	<b>Not Applicable</b>
	Test #2 All	Not applicable to DUT	<b>Not Applicable</b>
<b>Discussion:</b>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.12: TI MPI – Open	Test #1 PingPing and PingPong	Not applicable to DUT	<b>Not Applicable</b>
	Test #2 All	Not applicable to DUT	<b>Not Applicable</b>
<b>Discussion:</b>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.13: TI MPI – Mvapich1	Test #1 PingPing and PingPong	Not applicable to DUT	<b>Not Applicable</b>
	Test #2 All	Not applicable to DUT	<b>Not Applicable</b>
<b>Discussion:</b>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12.13: TI MPI – Mvapich2	Test #1 PingPing and PingPong	Not applicable to DUT	<b>Not Applicable</b>
	Test #2 All	Not applicable to DUT	<b>Not Applicable</b>
<b>Discussion:</b>			