



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

Michael Dulkan
 Flextronics (Israel) Ltd.
 1 Hastia str. Ramat Gavrial Ind. Zone Migdal Haemek
 Israel 23108
 Israel

August 21, 2009
 Report Rev1.3
 OFED Version: 1.4.1

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

Flextronics F-X430066 (8 Port SDR InfiniBand Switch)
Flextronics F-X430044 (24 Port DDR 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	Mandatory	PASS with Comments
10.2: IB Fabric Initialization	Mandatory	PASS
10.3: IB IPoIB Connected Mode	Mandatory	PASS
10.4: IB IPoIB Datagram Mode	Mandatory	PASS with Comments
10.6: IB SRP	Mandatory	PASS
12.1: TI iSER	Mandatory	Not Available
12.4: TI SDP	Mandatory	PASS


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

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

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 topology and device names
- v1.2 Fixed topology and device names. Added note about beta tests in results summary.
- V1.3 Changed topology diagram.

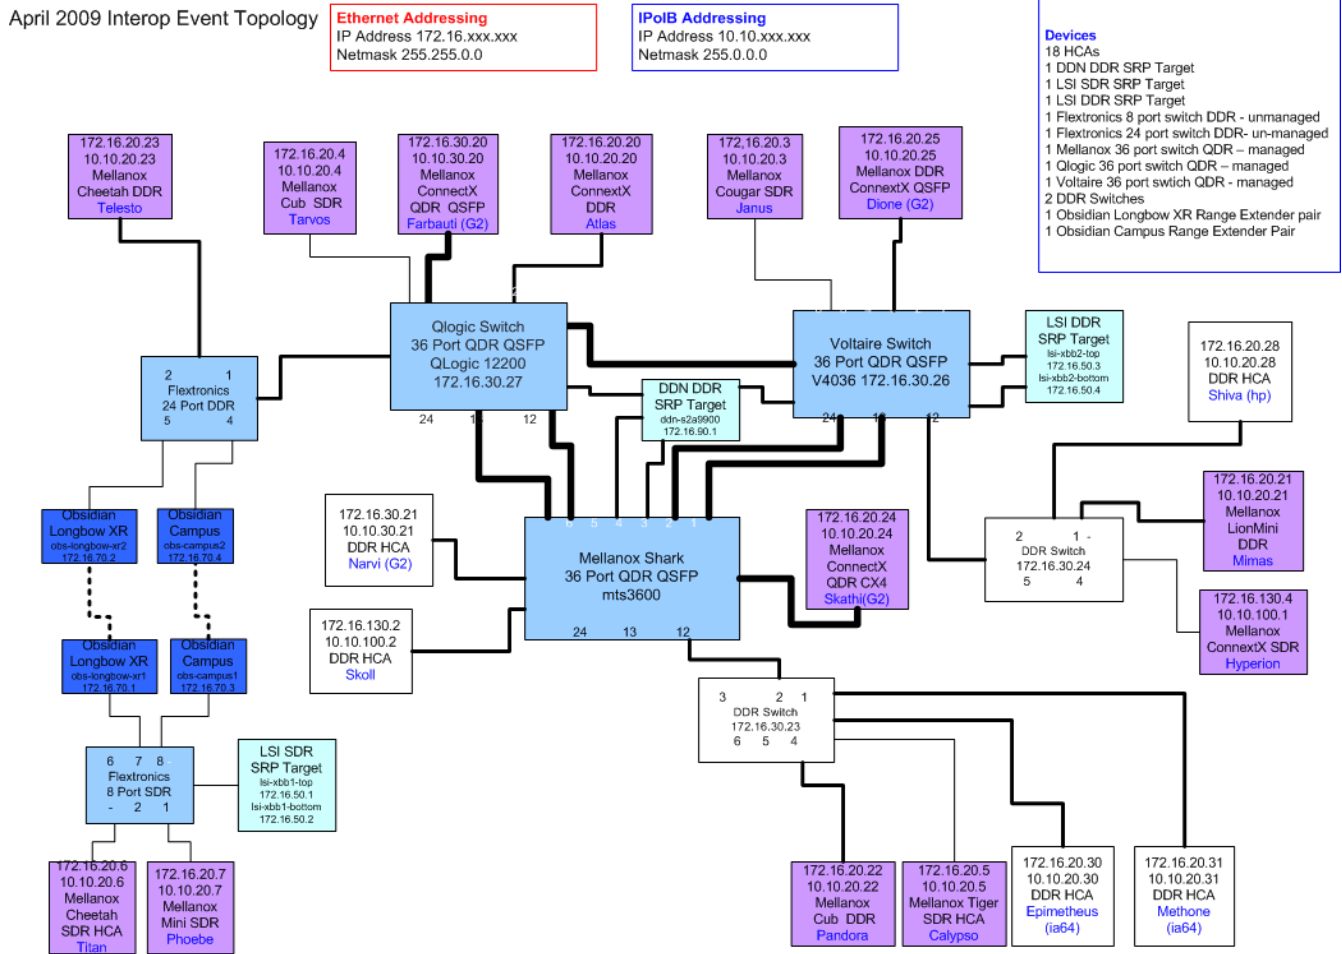
Table 2: Result Key

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, such as due to time limitations only a portion of the testing was performed.
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 specified parameters are valid at one extreme and invalid at the other.
Not Tested	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	Flextronics Ltd.	Firmware Rev	N/A
Model	F-X430066	Hardware Rev	A102/8
Speed	SDR 4x	IP Address in Fabric	N/A
Additional Comments/Notes			
Flextronics F-X430066 (8 Port SDR InfiniBand Switch)			

DUT #2 Details			
Manufacturer	Flextronics Ltd.	Firmware Rev	N/A
Model	F-X430044	Hardware Rev	I7
Speed	DDR 4x	IP Address in Fabric	N/A
Additional Comments/Notes			
Flextronics F-X430044 (24 Port DDR InfiniBand Switch)			

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	PASS with Comments
Discussion:			
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.			
This version of the test plan does not explicitly demand a link to be at the proper speed, the next version will. Logo testing in the 2 nd half of 2009 will require proper link speed between all link combinations. Narvi, Skoll, and the 12200 switch only link at SDR speeds even though they are rated to perform at speeds higher than SDR and the F-X430044 is a DDR switch.			

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

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	PASS
Discussion:			
<p>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.</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>			

<i>F-X430044</i>	All ports Armed/Active	No Dup GUIDs	No Port errors
OFED OpenSM (SM only)	PASS	PASS	PASS
QLogic 12200 (Managed Switch)	PASS	PASS	PASS
Mellanox MTS3600 (Managed Switch)	PASS	PASS	PASS
Voltaire v4036 (Managed Switch)	PASS	PASS	PASS

<i>F-X430066</i>	All ports Armed/Active	No Dup GUIDs	No Port errors
OFED OpenSM (SM only)	PASS	PASS	PASS
QLogic 12200 (Managed Switch)	PASS	PASS	PASS
Mellanox MTS3600 (Managed Switch)	PASS	PASS	PASS
Voltaire v4036 (Managed Switch)	PASS	PASS	PASS

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 10.3: IPoIB Connected Mode	Test #1	Automated Test Script	PASS
	Test #2	Automated Test Script	PASS
	Test #3	Automated Test Script	PASS
Discussion:			
<p>Test #1: 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>Test #2: 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>Test #3: 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	PASS with Comments
	Test #2	Automated Test Script	PASS
	Test #3	Automated Test Script	PASS
Discussion:			
<p>Test #1: 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.</p> <p>Test #2: 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>Test #3: 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	PASS
Discussion:			
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	Not Available
Discussion:			
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	PASS
	Test #2: SFTP	Automated Test Script	PASS
	Test #3: SCP	Automated Test Script	PASS
Discussion:			
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 #1-4	No problems detected	PASS
Discussion:			
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	Not Applicable
Discussion:			
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	Not Applicable
Discussion:			
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	Not Applicable
Discussion:			

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

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	Not Tested
Discussion:			
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	Not Tested
Discussion:			
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	Not Tested
Discussion:			
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	Not Tested
Discussion:			
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	Not Applicable
Discussion:			

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	Not Applicable
	Test #2 All	Not applicable to DUT	Not Applicable
Discussion:			

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	Not Applicable
	Test #2 All	Not applicable to DUT	Not Applicable
Discussion:			

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	Not Applicable
	Test #2 All	Not applicable to DUT	Not Applicable
Discussion:			

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	Not Applicable
	Test #2 All	Not applicable to DUT	Not Applicable
Discussion:			