



# OpenFabrics Alliance

## Interoperability Working Group (OFA-IWG)

### October 2007 Interoperability Event Report

UNH-IOL – 121 Technology Drive, Suite 2 – Durham, NH 03824 – +1-603-862-0090  
 OFILG – [ofalab@iol.unh.edu](mailto:ofalab@iol.unh.edu) – +1-603-862-5083

Amit Krig  
 Mellanox Technologies  
 2900 Stender Way  
 Santa Clara, CA 95054

November 9, 2007  
 Report Rev1.0

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

- Mellanox Technologies MHGH28-XTC HCA*
- Mellanox Technologies MHEA28-ITC HCA*
- Mellanox Technologies MHEA28-XT HCA*
- Mellanox Technologies MHGA28-ITC HCA*
- Mellanox Technologies MHGA28-XTC HCA*
- Mellanox Technologies MHES14-XT HCA*
- Mellanox Technologies MHES18-XTC HCA*
- Mellanox Technologies MHGS18-XTC HCA*

The test suite referenced in this report is available at the OFA website, at test time release 1.12 (September 12, 2007) was used:

<http://www.iol.unh.edu/services/testing/ofa/testplan.pdf>

<http://www.openfabrics.org/downloads/OFA-IWG%20Interoperability%20Test%20Plan-v1.12.pdf>

Test Procedure	IWG Test Status	Result/Notes
<a href="#">Group 01: IB Link Initialize</a>	<b>Mandatory</b>	Passed – no issues seen
<a href="#">Group 02: IB Fabric Initialization</a>	<b>Mandatory</b>	Passed – no issues seen
<a href="#">Group 03: IB IPoIB</a>	<b>Mandatory</b>	Passed – no issues seen
<a href="#">Group 04: TI iSER</a>	Beta	<b>Refer to Comments</b>
<a href="#">Group 05: IB SRP</a>	<b>Mandatory</b>	Passed – no issues seen
<a href="#">Group 06: TI SDP</a>	<b>Mandatory</b>	Passed – no issues seen
<a href="#">Group 07: IB SM Failover</a>	Beta	Not Tested
<a href="#">Group 08: TI MPI - OSU</a>	Beta	Not Tested
<a href="#">Group 09: TI MPI - Intel</a>	Beta	Passed – no issues seen
<a href="#">Group 10: TI uDAPL</a>	Beta	<b>Refer to Comments</b>
<a href="#">Group 12: FibreChannel Gateway(IB)</a>	Beta	Not applicable to DUT
<a href="#">Group 13: Ethernet Gateway(IB)</a>	Beta	Not applicable to DUT

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

Testing Completed 10/30/2007

Bob Noseworthy  
<mailto:ren@iol.unh.edu>



Review Completed 11/29/2007

Mikkel Hagen  
<mailto:mhagen@iol.unh.edu>

## 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/>

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: A303 D24B 3F7D 0E0D 27F2 B8BC 5FA0 1FC6

SHA-1 Fingerprint: 7BD1 A2EE 89DC AB98 2E32 F36A A9E6 E865 A0EE 88EE

## Report Revision History

v1.0 Initial Release


## Table 1: 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 2: 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 Technologies	Firmware Rev	N/A
Model	MHGH28-XTC	Hardware Rev	N/A
Speed	DDR	IP Address in Fabric	132.177.126.225 / 10.1.40.225
Additional Comments/Notes			
AKA "ConnectX DDR" - Dual-port, 20Gb/s w/ media adapter support			

DUT #2 Details			
Manufacturer	Mellanox Technologies	Firmware Rev	4.8.2
Model	MHEA28-1TC	Hardware Rev	A1
Speed	SDR	Driver Version	OFED 1.2 RC1
		IP Address in Fabric	132.177.126.228 / 10.1.40.228
Additional Comments/Notes			
AKA "Lion Cub HCA" - Dual-port, 10Gb/s w/ media adapter support, PCIe x8, 128MB memory, tall bracket, RoHS R5			

<b>DUT #3 Details</b>			
Manufacturer	Mellanox Technologies	Firmware Rev	5.2.0
Model	MHEA28-XT	Hardware Rev	A2
Speed	SDR	Driver Version	OFED 1.2 RC1
		IP Address in Fabric	132.177.126.200 / 10.1.40.200
<b>Additional Comments/Notes</b>			
AKA “Lion Mini HCA” - Dual-port, 10Gb/s w/ media adapter support, PCIe x8, mem-free, tall bracket			

<b>DUT #4 Details</b>			
Manufacturer	Mellanox Technologies	Firmware Rev	4.8.2
Model	MHGA28-ITC	Hardware Rev	A2
Speed	DDR	Driver Version	OFED 1.2 RC1
		IP Address in Fabric	132.177.126.205 / 10.1.40.205
<b>Additional Comments/Notes</b>			
AKA “Lion Cub DDR” - Dual-port, 20Gb/s w/ media adapter support, PCIe x8, 128MB memory, tall bracket, RoHS R5			

<b>DUT #5 Details</b>			
Manufacturer	Mellanox Technologies	Firmware Rev	5.2.0
Model	MHGA28-XTC	Hardware Rev	A2
Speed	DDR	Driver Version	OFED 1.2 RC1
		IP Address in Fabric	132.177.126.221 / 10.1.40.221
<b>Additional Comments/Notes</b>			
AKA “Lion Mini DDR” - Dual-port, 20Gb/s w/ media adapter support, PCIe x8, mem-free, tall bracket, RoHS R5			

<b>DUT #6 Details</b>			
Manufacturer	Mellanox Technologies	Firmware Rev	1.2.0
Model	MHES14-XT	Hardware Rev	A4
Speed	SDR	Driver Version	OFED 1.2 RC1
		IP Address in Fabric	132.177.126.204 / 10.1.40.204
<b>Additional Comments/Notes</b>			
AKA “Tiger HCA” - Single-port, 10Gb/s w/ media adapter support, PCIe x4, mem-free, tall bracket			

<b>DUT #7 Details</b>			
Manufacturer	Mellanox Technologies	Firmware Rev	1.2.0
Model	MHES18-XTC	Hardware Rev	A2
Speed	SDR	Driver Version	OFED 1.2 RC1
		IP Address in Fabric	132.177.126.232 / 10.1.40.232
<b>Additional Comments/Notes</b>			
AKA “Cheetah HCA” - Single-port, 10Gb/s w/ media adapter support, PCIe x8, mem-free, tall bracket, RoHS R5			

<b>DUT #8 Details</b>			
Manufacturer	Mellanox Technologies	Firmware Rev	1.2.0
Model	MHGS18-XTC	Hardware Rev	A3
Speed	DDR	Driver Version	OFED 1.2 RC1
		IP Address in Fabric	132.177.126.229 / 10.1.40.229
<b>Additional Comments/Notes</b>			
AKA “Cheetah DDR HCA” - Single-port, 20Gb/s w/ media adapter support, PCIe x8, mem-free, tall bracket, RoHS R5			

## Mandatory Tests - IB Device Test Summary Results:

The following tables detail results for tests identified by the OFA-IWG as mandatory tests for the OFA Interoperability Logo Program (OFILP) per the OFA-IWG Interoperability Test Plan Release 1.12 (September 12, 2007)

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 1: IB Link Initialize	Test #1: Phy link up all ports	Link partners link as expected	<b>PASS</b>
	Test #2: Logical link up all ports switch SM	ibstatus reports active links	<b>PASS</b>
	Test #3: Logical link up all ports HCA SM	ibstatus reports active links	<b>PASS</b>

### Discussion: Test #1: Phy link up all ports

DDR cables were used for all link tests. Device LEDs validated visually. Link width and link speed validated via use of “ibdiagnet -lw 4x” and ibdiagnet -ls 10”

### Discussion: Test #2 & 3: Logical link up all ports switch SM / HCA SM

The switch-under-test's SM, OFED's OpenSM, and the Cisco High performance SM (HSM) were all used to validate that the link could be brought to the Active state as verified via the “ibstatus” command. All SMs were initially off as the switch was powered up. After checking for an initial physical link, the SM under test was started and the state of the link verified. Note, for HCA to HCA link checks, the nature of the direct cable connection prevented the validation of any SM but the OpenSM. Refer to the table below for specific link configurations tested.

For Mellanox HCAs	Switch SM	OpenSM	Cisco High perf. SM
Cisco SFS 7000D 24-port DDR (Managed Switch)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Cisco SFS 7000P 24-port SDR (Managed Switch)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Cisco SFS 7008 96-port SDR (Managed Switch)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Flextronics 8 Modular SDR (Unmanaged Switch)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Flextronics 24 Modular DDR (Unmanaged Switch)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Mellanox MHEA28-1TC “Lion Cub SDR” (HCA)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Mellanox MHEA28-XT “Lion Mini SDR” (HCA)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Mellanox MHGA28-1TC “Lion Cub DDR” (HCA)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Mellanox MHGA28-XTC “Lion Mini DDR” (HCA)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Mellanox MHES14-XT “Tiger SDR” (HCA)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Mellanox MHES18-XTC “Cheetah SDR” (HCA)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Mellanox MHGS18-XTC “Cheetah DDR” (HCA)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Mellanox MHGH28-XTC “ConnectX DDR” (HCA)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Qlogic SilverStorm 9024 24-port DDR (Managed Switch)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Qlogic SilverStorm 9040 40-port DDR (Managed Switch)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Qlogic QHT-7140 {SDR PCIe x8} (HCA)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Qlogic QLE-7140 {SDR PCIe x8} (HCA)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Qlogic QLE-7240 {DDR PCIe x8} (HCA)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Qlogic QLE-7280 {DDR PCIe x16} (HCA)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Voltaire IPSTOR (iSER Target / FC Gateway)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
<b>Group 2: IB Fabric Initialization</b>	Test #1: Verify all SMs configure fabric	Port is Active with all SMs	<b>PASS</b>
<b>Discussion: Test #1: Verify all SMs configure fabric</b>			
The fabric configuration shown in Figure 1 was used for this test. 'ibdiagnet -c 1000' showed no Port errors 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 on each host 'ibstatus' to validate speed, width and link state. Refer to the table below for SM details.			

For Mellanox HCAs	All ports Armed/Active	No Dup GUIDs	No Port errors
Cisco High performance SM(HSM) Host (HCA)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Cisco SFS 7000D 24-port DDR (Switch)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Cisco SFS 7000P 24-port SDR (Switch)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Cisco SFS 7008 96-port SDR (Switch)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
OFED 1.2.5.1 OpenSM	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Qlogic SilverStorm 9024 24-port DDR (Switch)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>
Qlogic SilverStorm 9040 40-port DDR (Switch)	<b>PASS</b>	<b>PASS</b>	<b>PASS</b>

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
<b>Group 3: IPoIB</b>	Test #1: Ping Test all to all	HCAs can ping all to all with all byte sizes	<b>PASS</b>
	Test #2: Connect disconnect host	Connectivity functions after topology change	<b>PASS</b>
	Test #3: File Transfer Procedure	4MB file SFTP transfers were validated	<b>PASS</b>
<b>Discussion: Test #1: Ping all to all</b>			
IP connectivity to all HCAs using each class of SMs was validated. SMs include Cisco switch SMs, Cisco HSM, OpenSM, Qlogic switch SMs. Note that to achieve ICMP echo requests (pings) with a packet size of 64, 256, 511, 512, 1024, 1025, 2044, 4096, 8192, 16384, 32768, 65507, ping -s # was utilized where # corresponds respectively in the following list: 36, 228, 483, 484, 996, 997, 2016, 4068, 8164, 16356, 32740, 65479 Note: One non-Mellanox HCA was observed to not support pings beyond 2048 in size.			
<b>Discussion: Test #2: Disconnect and Reconnect HCA</b>			
Each class of SMs was tested (Cisco switch SMs, Cisco HSM, OpenSM, Qlogic switch SMs). Refer to the default fabric configuration. All HCAs were connected except the Qlogic x16 DDR HCA (10.1.60.240). All SMs were disabled, all switches were then power cycled. The SM under test was enabled, and all HCAs were confirmed to be reachable except for the station at 10.1.40.200. The Mellanox ConnectX HCA (10.1.40.225) was removed from the fabric. All HCAs were confirmed to be reachable except the two disconnected HCAs. The Qlogic x16 DDR HCA was then connected to the fabric. All HCAs are now reachable except the ConnectX. The Mellanox ConnectX is now connected to the Qlogic 9040 Silverstorm DDR switch rather than the Qlogic 9024 Silverstorm DDR switch. All HCAs are now reachable. Using host system 10.1.40.225, "ifconfig ib0 down" results in a loss of IP connectivity, "ifconfig ib0 up" results in a restoration of IP connectivity			
<b>Discussion: Test #3: File transfer procedure</b>			
For expediency, SCP was utilized as FTP servers were not configured at test time. File transfer was validated from the Mini-SDR host (10.1.40.225) to each other station by sending a 4MB file to the remote station and then copying it back and comparing the file. Tests were repeated 4 times. All file transfers completed successfully and the file was identical in all cases.			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 5: IB SRP	Test #1-4	No issues observed	<b>PASS</b>
<b>Discussion: Test #1-4</b>			
<p>Note, the test plan v1.12 was modified from requiring only “2 HCAs” to include representative hosts from all vendors used for validation. The LSI Engenio 6498 was utilized as the SRP target for the HCAs and switches under test. For this test, hosts logged into the SRP target and performed reads via the command 'dd if=/dev/sdb of=/dev/null count=6 bs=10M'. Note, during testing the SM should not be stopped until login and data transfer on all connections is confirmed, thus ensuring all devices have properly logged in and commenced transfers before terminating the current SM. The host was observed to complete the current in-progress transfer when the SM is killed. When disconnected and reconnected, a host could still perform a dd operation</p>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 6: TI SDP	Test #1: netperf procedure	All tests completed successfully	<b>PASS</b>
	Test #2: FTP procedure	4MB transfers were validated	<b>PASS</b>
	Test #3: IB SCP Procedure	SCP file transfer validated	<b>PASS</b>
<b>Discussion: Test #1: netperf procedure</b>			
<p>Module ib_sdp is loaded. Netperf is run to and from all hosts in the fabric. Note message sizes 10, 100, 1000, 10000 and buffer sizes of 1024 and 6000 were validated on all connections. All netperf operations were observed to complete successfully.</p>			
<b>Discussion: Test #2: FTP procedure</b>			
<p>For expedience, SFTP servers were configured and a 4MB file was transferred to and from each host and a binary comparison performed to validate the transfer. 'lsmod   grep sdp' was used to validate that ib_sdp was loaded with &gt;0 dependencies.</p>			
<b>Discussion: Test #3: IB SCP Procedure</b>			
<p>scp was used successfully between all hosts to transfer files while ib_sdp was loaded.</p>			

## Beta Tests - IB Device Test Results:

The following table details results for tests identified by the OFA-IWG as beta tests for the OFA Interoperability Logo Program (OFILP) per the OFA-IWG Interoperability Test Plan Release 1.12 (September 12, 2007)

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 4: TI iSER	Test #1-5	Issues observed	<b>Refer to Comments</b>
<b>Discussion: Test Results</b>			
<p>When iSER operations are performed from 11 HCAs sequentially, no issues are observed.  When the same HCAs perform simultaneous iSER write operations to different locations on the target disk, issues were observed. At this time the issue is believed to be related to the iSER target and not the IB fabric.</p> <p>The current test plan (v1.12) makes no distinction of how the dd operations should be performed, stating simply “Run basic dd application from iSER initiator host connected to target”. Nor is there any indication of how multiple HCAs should be investigated (ex: test all HCAs sequentially, simultaneously, etc). With additional slight clarification in the test plan, this Beta procedure could advance to Mandatory in future Interoperability Events.</p>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 7: IB SM Failover/Handover	Test #1-4	Not tested	<b>Not Tested</b>
<b>Discussion: Test Results</b>			
These tests were not performed during the October 2007 Interoperability Event			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 8: TI MPI – Ohio State University	Test #1-14	Not tested	<b>Not Tested</b>
<b>Discussion: Test Results</b>			
These tests were not performed during the October 2007 Interoperability Event			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 9: MPI – Intel MPI	Test #1-21	No issues observed	<b>PASS</b>
<b>Discussion: Test Results</b>			
<p>A cluster of 8 nodes consisting of all Mellanox HCAs (one of each DUT type under test) was constructed.  Intel MPI was configured on the cluster and Intel MPI Benchmark tests were performed.  All mpich2-tests and IntelMPITESTs were completed with passing results.  Result logs are available upon request.</p>			



Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 10: TI uDAPLTEST Commands	Test #1-10	Issues observed	Refer to Comments
<b>Discussion: Test Results</b>			
<p>Time was taken to test and retest throughout October up till October 30, 2007 including re-investigations of this portion of the test plan to investigate issues seen during testing.</p> <p>Tests during the Interoperability Event were run utilizing the test script published in v1.12 of the test plan.</p> <p>The following failure mode was observed when the script was executed. This failure was not observed consistently.</p> <p>Test 3.1 would complete successfully but report the following error message:</p> <pre>Test Error: Send-reaping DTO problem, status = FAILURE dapltest: dapl/common/dapl_ep_free.c:114: dapl_ep_free: Assertion `ep_ptr-&gt;param.ep_state == DAT_EP_STATE_DISCONNECTED    ep_ptr-&gt;param.ep_state == DAT_EP_STATE_UNCONNECTED' failed. Test Error: Send-reaping DTO problem, status = FAILURE Test Error: Recv-reaping DTO problem, status = FAILURE</pre> <p>Subsequently, tests 3.2, 3.3 and 3.4 would fail with the following error message:</p> <pre>Warning: conn_event_wait DAT_CONNECTION_EVENT_NON_PEER_REJECTED DT_cs_Client: bad connection event</pre> <p>In at least one test run, the Mellanox MHES14-XT “Tiger SDR” (HCA) was common to all test scenarios where the above test error was observed. Insufficient data currently exists to further isolate the root cause of the observed issue.</p> <p>Due to these pending issues, further testing is required for this Beta test process.</p>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 11: iWARP Connectivity	Test #1-8	Not applicable to DUT	Not Applicable
<b>Discussion: Test Results</b>			
The OFA Logo Program does not require these tests for an HCA under test.			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 12: Fibre Channel Gateway (IB)	Test #1-10	Not applicable to DUT	Not Applicable
<b>Discussion: Test Results</b>			
The OFA Logo Program does not require these tests for an HCA under test.			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 13: Ethernet Gateway (IB)	Test #1-7	Not applicable to DUT	Not Applicable
<b>Discussion: Test Results</b>			
The OFA Logo Program does not require these tests for an HCA under test.			