



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 – +1-603-862-5083

Clive Hall
 Cisco Systems
 170 West Tasman Drive
 San Jose, CA 95134

November 9, 2007
 Report Rev1.0

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

Cisco Systems SFS 7000D 24-port DDR Switch
Cisco Systems SFS 7000P 24-port SDR Switch
Cisco Systems SFS 7008 96-port SDR Switch

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
Group 01: IB Link Initialize	Mandatory	Passed – no issues seen
Group 02: IB Fabric Initialization	Mandatory	Passed with Comments
Group 03: IB IPoIB	Mandatory	Passed – no issues seen
Group 04: TI iSER	Beta	Refer to Comments
Group 05: IB SRP	Mandatory	Passed – no issues seen
Group 06: TI SDP	Mandatory	Passed – no issues seen
Group 07: IB SM Failover	Beta	Not Tested
Group 08: TI MPI - OSU	Beta	Not applicable to DUT
Group 09: TI MPI - Intel	Beta	Not applicable to DUT
Group 10: TI uDAPL	Beta	Not applicable to DUT
Group 12: FibreChannel Gateway(IB)	Beta	Not applicable to DUT
Group 13: Ethernet Gateway(IB)	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

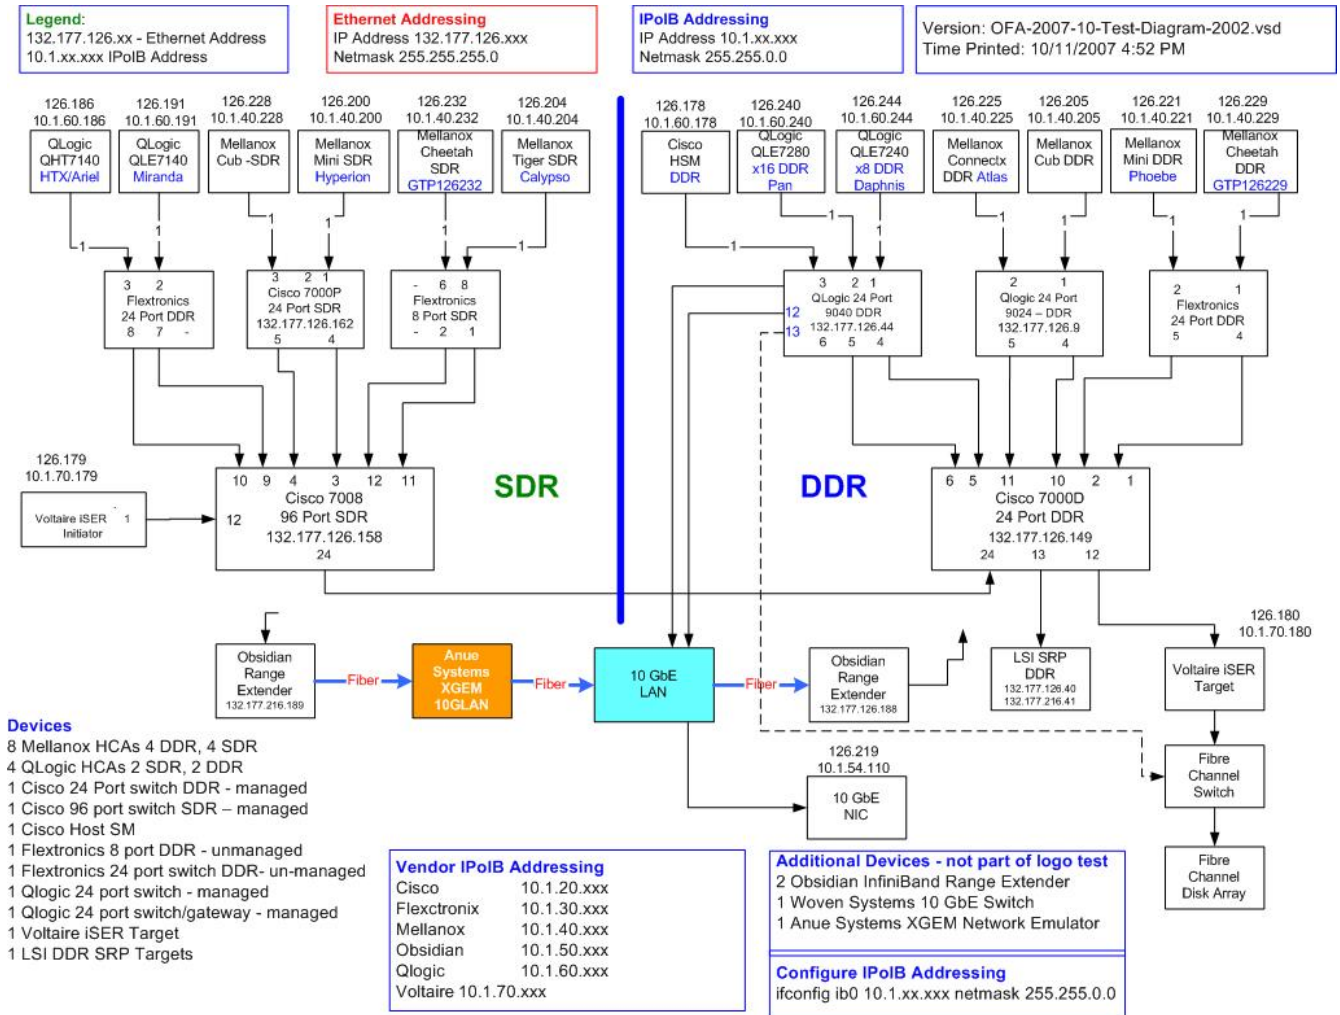
Table 1: Result Key

The following table contains possible results and their meanings:

Result:	Description:
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 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	Cisco Systems, Inc	Firmware Rev	2.9.0
Model	SFS 7000D	Hardware Rev	02
Speed	DDR	IP Address in Fabric	132.177.126.149
Additional Comments/Notes			
Managed 24-port DDR Switch			

DUT #2 Details			
Manufacturer	Cisco Systems, Inc	Firmware Rev	2.9.0
Model	SFS 7000P	Hardware Rev	N/A
Speed	SDR	IP Address in Fabric	132.177.126.162
Additional Comments/Notes			
Managed 24-port SDR Switch			

DUT #3 Details			
Manufacturer	Cisco Systems, Inc	Firmware Rev	2.9.0
Model	SFS 7008	Hardware Rev	C0
Speed	SDR	IP Address in Fabric	132.177.126.158
Additional Comments/Notes			
Managed 96-port SDR Switch			

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	PASS
	Test #2: Logical link up all ports switch SM	ibstatus reports active links	PASS
	Test #3: Logical link up all ports HCA SM	ibstatus reports active links	PASS

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 Cisco SFS7008 SFS7000P & SFS7000D Switches	Switch SM	OpenSM	Cisco High perf. SM
Cisco SFS 7000D 24-port DDR (Managed Switch)	PASS	PASS	PASS
Cisco SFS 7000P 24-port SDR (Managed Switch)	PASS	PASS	PASS
Cisco SFS 7008 96-port SDR (Managed Switch)	PASS	PASS	PASS
Flextronics 8 Modular SDR (Unmanaged Switch)	PASS	PASS	PASS
Flextronics 24 Modular DDR (Unmanaged Switch)	PASS	PASS	PASS
Mellanox MHEA28-1TC “Lion Cub SDR” (HCA)	PASS	PASS	PASS
Mellanox MHEA28-XT “Lion Mini SDR” (HCA)	PASS	PASS	PASS
Mellanox MHGA28-1TC “Lion Cub DDR” (HCA)	PASS	PASS	PASS
Mellanox MHGA28-XTC “Lion Mini DDR” (HCA)	PASS	PASS	PASS
Mellanox MHES14-XT “Tiger SDR” (HCA)	PASS	PASS	PASS
Mellanox MHES18-XTC “Cheetah SDR” (HCA)	PASS	PASS	PASS
Mellanox MHGS18-XTC “Cheetah DDR” (HCA)	PASS	PASS	PASS
Mellanox MHGH28-XTC “ConnectX DDR” (HCA)	PASS	PASS	PASS
Qlogic SilverStorm 9024 24-port DDR (Managed Switch)	PASS	PASS	PASS
Qlogic SilverStorm 9040 40-port DDR (Managed Switch)	PASS	PASS	PASS
Qlogic QHT-7140 {SDR PCIe x8} (HCA)	PASS	PASS	PASS
Qlogic QLE-7140 {SDR PCIe x8} (HCA)	PASS	PASS	PASS
Qlogic QLE-7240 {DDR PCIe x8} (HCA)	PASS	PASS	PASS
Qlogic QLE-7280 {DDR PCIe x16} (HCA)	PASS	PASS	PASS
Voltaire IPSTOR (iSER Target / FC Gateway)	PASS	PASS	PASS

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 2: IB Fabric Initialization	Test #1: Verify all SMs configure fabric	Port is Active with all SMs	PASS

Discussion: Test #1: Verify all SMs configure fabric

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.

Note: During testing, the OpenSM 1.2.5.1 was observed to crash on several occasions immediately following the OpenSM host HCA being connected directly to the SFS7000D. At this time, no information indicates an issue with the SFS7000D and instead an issue with the OpenSM 1.2.5.1 is suspected. The OpenSM could successfully be restarted without incident once the connection is in place.

For Cisco SFS7008 SFS7000P & SFS7000D Switches	All ports Armed/Active	No Dup GUIDs	No Port errors
Cisco High performance SM(HSM) Host (HCA)	PASS	PASS	PASS
Cisco SFS 7000D 24-port DDR (Switch)	PASS	PASS	PASS
Cisco SFS 7000P 24-port SDR (Switch)	PASS	PASS	PASS
Cisco SFS 7008 96-port SDR (Switch)	PASS	PASS	PASS
OFED 1.2.5.1 OpenSM	PASS	PASS	PASS
Qlogic SilverStorm 9024 24-port DDR (Switch)	PASS	PASS	PASS
Qlogic SilverStorm 9040 40-port DDR (Switch)	PASS	PASS	PASS

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 3: IPoIB	Test #1: Ping Test all to all	HCA's can ping all to all with all byte sizes	PASS
	Test #2: Connect disconnect host	Connectivity functions after topology change	PASS
	Test #3: File Transfer Procedure	4MB file SFTP transfers were validated	PASS
Discussion: Test #1: Ping all to all			
<p>IP connectivity to all HCA's using each class of SMs was validated. SMs include Cisco switch SMs, Cisco HSM, OpenSM, Qlogic switch SMs.</p> <p>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</p> <p>Note: One non-Cisco device was observed to not support pings beyond 2048 in size.</p>			
Discussion: Test #2: Disconnect and Reconnect HCA			
<p>Each class of SMs was tested (Cisco switch SMs, Cisco HSM, OpenSM, Qlogic switch SMs).</p> <p>Refer to the default fabric configuration. All HCA's were connected except the Qlogic x16 DDR HCA (10.1.60.240). All SMs were disabled, all switches were then power cycled.</p> <p>The SM under test was enabled, and all HCA's 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 HCA's were confirmed to be reachable except the two disconnected HCA's.</p> <p>The Qlogic x16 DDR HCA was then connected to the fabric. All HCA's 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 HCA's are now reachable.</p> <p>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</p>			
Discussion: Test #3: File transfer procedure			
<p>For expediency, SCP was utilized as FTP servers were not configured at test time.</p> <p>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.</p> <p>Tests were repeated 4 times. All file transfers completed successfully and the file was identical in all cases.</p>			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 5: IB SRP	Test #1-4	No issues observed	PASS
Discussion: Test #1-4			
<p>Note, the test plan v1.12 was modified from requiring only "2 HCA's" to include representative hosts from all vendors used for validation. The LSI Engenio 6498 was utilized as the SRP target for the HCA's 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.</p> <p>The host was observed to complete the current in-progress transfer when the SM is killed.</p> <p>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	PASS
	Test #2: FTP procedure	4MB transfers were validated	PASS
	Test #3: IB SCP Procedure	SCP file transfer validated	PASS
Discussion: Test #1: netperf procedure			
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.			
Discussion: Test #2: FTP procedure			
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 >0 dependencies.			
Discussion: Test #3: IB SCP Procedure			
scp was used successfully between all hosts to transfer files while ib_sdp was loaded.			

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	Refer to Comments
Discussion: Test Results			
<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	Not Tested
Discussion: Test Results			
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 applicable to DUT	Not Applicable
Discussion: Test Results			
The OFA Logo Program does not require these tests for a IB Fabric switch under test.			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 9: MPI – Intel MPI	Test #1-21	Not applicable to DUT	Not Applicable
Discussion: Test Results			
The OFA Logo Program does not require these tests for a IB Fabric switch under test.			

Test Number and Name	Part(s)	Summary Note(s)	Result(s)
Group 10: TI uDAPLTEST Commands	Test #1-10	Not applicable to DUT	Not Applicable
Discussion: Test Results			
The OFA Logo Program does not require these tests for a IB Fabric switch under test.			

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