OpenFabrics Alliance

Interoperability Logo Group (OFILG)

January 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

Clifford Cole
Intel Corporation
780 Fifth Avenue
Suite 140
King of Prussia, PA 19406

Date:	16 June 2015
Report Revision:	2.0
OFED Version on Compute Nodes:	3.12-1
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):

Intel QLE7340 *Intel QLE7342*

The test suite referenced in this report is available at the UNH-IOL website. Release 1.51 (2014-Sep-23) was used.

https://iol.unh.edu/ofatestplan

The following table highlights the Mandatory test required for the OpenFabrics Interoperability Logo for the InfiniBand HCA device class per the Test Plan and 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 with Comments
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 with Comments
13.4: TI uDAPL	Mandatory	PASS
13.5: TI RDMA Basic Interoperability	Mandatory	PASS
13.6: TI RDMA Stress	Mandatory	PASS
13.7: RSockets	Mandatory	PASS with Comments
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 June 16, 2015

Dave Wyman
dwyman@iol.unh.edu

Digitally signed by UNH-IOL
Date: 2015.07.22 09:25:51 -04'00'

Reviewed & Issued July 22, 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 HCA).

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 with Comments
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 with Comments
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 with Comments
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:

https://www.iol.unh.edu/testing/reports/certificate-install

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
- v2.0 updated IPoIB, NFSoRDMA and RSockets tests

Configuration Files

Description	Attachment
Scientific Linux 7.0 Configuration File	
OFED 3.12-1 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.

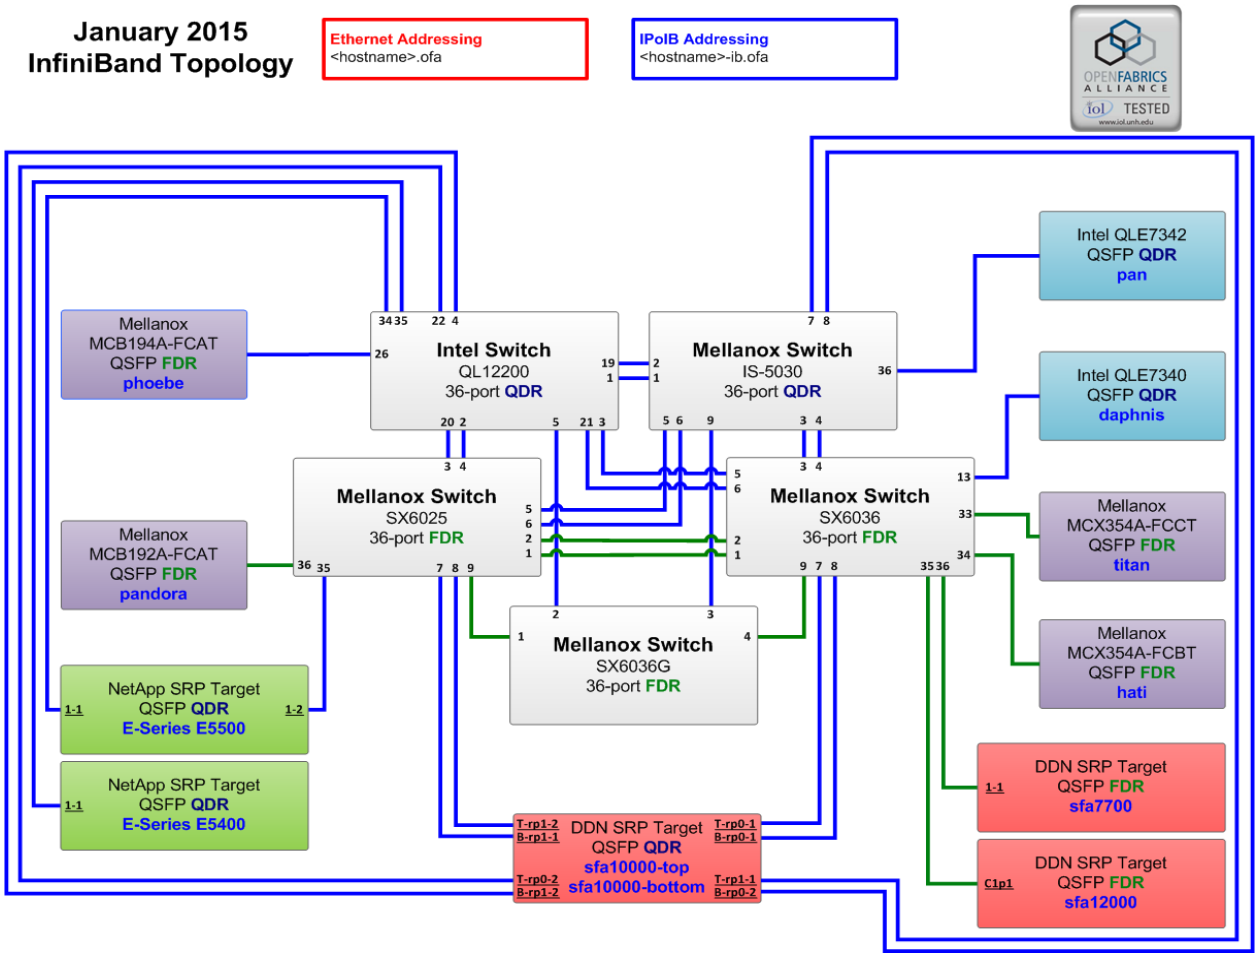

DUT #1 Details			
Manufacturer:	Intel	Firmware Revision:	N/A *
Model:	QLE7340	Hardware Revision:	2
Speed:	QDR	Located in Host:	daphnis
Firmware MD5sum:	N/A		
Additional Comments / Notes:			
* Contained in OFED 3.12-1			

DUT #2 Details			
Manufacturer:	Intel	Firmware Revision:	N/A *
Model:	QLE7342	Hardware Revision:	2
Speed:	QDR	Located in Host:	pan
Firmware MD5sum:	N/A		
Additional Comments / Notes:			
* Contained in OFED 3.12-1			

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		QLE7340	QLE7342
Intel 12200 (Switch) – QDR		PASS	PASS
Mellanox SX6025 (Switch) – FDR		PASS	PASS
Mellanox SX6036 (Switch) – FDR		PASS	PASS
Mellanox IS-5030 (Switch) – QDR		PASS	PASS
Mellanox SX6036G (Switch) – FDR		PASS	PASS
DataDirect Networks SFA12000 (SRP Target) – FDR		PASS	PASS
DataDirect Networks SFA10000 (SRP Target) – QDR		PASS	PASS
DataDirect Networks SFA7700 (SRP Target) – FDR		PASS	PASS
NetApp Soyuz (SRP Target) – QDR		PASS	PASS
NetApp Pikes Peak (SRP Target) – QDR		PASS	PASS
Host: hati	HCA: MCX354A-FCBT (FDR)	PASS	PASS
Host: titan	HCA: MCX354A-FCCT (FDR)	PASS	PASS
Host: phoebe	HCA: MCB194A-FCAT (FDR)	PASS	PASS
Host: pandora	HCA: MCB192A-FCAT (FDR)	PASS	PASS
Host: pan	HCA: QLE7342 (QDR)	PASS	NA
Host: daphnis	HCA: QLE7340 (QDR)	NA	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 with Comments**	**PASS**	**PASS**
Result Discussion:			
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 65493 or higher. To reproduce the issue, follow the testplan IPoIB section 11.3.4 on page 48. The problem will occur upon execution of #5 in Step A, just after the paired device's InfiniBand interface has been removed from the ARP table. If the command run in #5, Step A, 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. 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.			

11.4: IPoIB Datagram Mode

Subnet Manager	Part A	Part B	Part C
OpenSM	**PASS with Comments**	**PASS**	**PASS**
Result Discussion:			
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 11.4.4 on page 51. The problem will occur upon execution of #5 in Step A, just after the paired device's InfiniBand interface has been removed from the ARP table. If the command run in #5, Step A, 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. 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.			

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:	
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	**PASS with Comments**
Result Discussion:	
No DUTs were able to complete all 4 sections of the Connectathon test suite in this Logo event. A subset of devices were able to complete the Basic, Locking, and Special sections of the suite. The General section exits with an error which may not be caused by the vendor hardware. To reproduce the issues refer to the ofatestplan NFSoRDMA Test Procedure in section 13.2.2 on page 72. Similar issues were observed with other pairs of devices with no Intel device in the pair. Connectathon suite is no longer maintained at connectathon.org. Further analysis reveals the failure occurs due to an incompatibility with the General section of the Connectathon suite and new kernel versions. In March 2015, IOL obtained a version of Connectathon renamed cthon04 from a github repository maintained by Steve Dickson. The repository is no longer available there but is now found on https://fedorapeople.org/cgit/steved/public_git/cthon04.git/. This git repo has not been tested since it was pulled down from github so the repository at the provided link may have progressed since the time of testing. However, the version previously obtained can be requested from IOL to reproduce the passing results. The workaround is to use the Basic, Locking, and Special sections from the Connectathon used in previous Logo events (and also available on demand from IOL), substituting in the General section from the cthon04 version. This workaround results in a passing result for all tests for each adapter acting as both client and server with the exception of XXXX and YYYY HCAs. XXXX and YYYY were unable to successfully complete the modified Connectathon test suite while running as the server but were successfully able to complete as the client. This result was also observed in the January and May 2014 Logo Event Reports and is not a fault of the Intel DUTs.	

13.4: TI uDAPL

Subnet Manager	Result
OpenSM	**PASS**
Result Discussion:	
All communications using DAPL 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:	
DUTs were noted to hang indefinitely when acting as either server or client during initial testing. When connected to another Intel device, each device was unable to complete RDMA Asynchronous, Blocking, and Non-blocking tests and instead hung on a polling thread. When connected to a non-Intel device the behavior is the same. Note that when two non-Intel devices are connected, the test performs properly. The issue was also observed when the Intel device was directly connected to an HCA. Logs for these issues are available from the UNH-IOL OFILG. Further analysis by rsockets maintainer Sean Hefty revealed an incompatibility between qib0 driver and librdmacm_1.0.19 included with OFED-3.12-1. After removing librdmacm_1.0.19 and installing maintainer provided librdmacm_1.0.21 the problem is resolved. The updated librdmacm will be included with OFED-3.18 resolving this issue for future releases.	

13.8: TI MPI – Open

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