



OpenFabrics Alliance

Interoperability Working Group (OFA-IWG)

September 2010 Interoperability GA Event Report

UNH-IOL – 121 Technology Drive, Suite 2 – Durham, NH 03824 – +1-603-862-0090
 OFILG – ofalab@iol.unh.edu

Jim Ryan
 Intel Corporation
 1501 South Mopac Suite 400
 Austin, TX 78746

January 10, 2011
 Report Revision 2.0
 OFED Version 1.5.2

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

Intel-NE NE0200101E3CX0104 Single Port CX4 RNIC

The test suite used for the tests described in this report is *OFA-IWG Interoperability Test Plan Release 1.35* (July 27, 2010 DRAFT). It is available at the OFA Web site:

http://www.iol.unh.edu/services/testing/ofa/testsuites/OFA-IWG_Interoperability_Test_Plan-v1.35.pdf

The OFILG gratefully acknowledges the **iWARP Consortium** for developing and contributing the testing software used for test 11.5 iWARP Connectivity.

Results Overview - Mandatory Tests

Test Procedure	IWG Test Status	Result/Notes
11.1 Ethernet Link Initialization	Mandatory	PASS with Comments
11.2 Ethernet Fabric Initialization	Mandatory	Not Available
11.5 iWARP Connectivity	Mandatory	PASS with Comments
12.11 Open MPI - homogeneous	Mandatory	PASS
12.12 MVAPICH2 - homogeneous	Mandatory	PASS

Testing Completed 2010-Oct-1

Jeff Laird
 laird@iol.unh.edu

Review Completed 2011-Jan-10

Bob Noseworthy
 ren@iol.unh.edu

Results Overview - All Tests

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, or with the test plan procedures.

Test Procedure	IWG Test Status	Result/Notes
11.1 Ethernet Link Initialization	Mandatory	PASS with Comments
11.2 Ethernet Fabric Initialization	Mandatory	Not Available
11.3 Ethernet Fabric Reconvergence	Beta	Not Available
11.4 Ethernet Fabric Failover	Beta	Not Available
11.5 iWARP Connectivity	Mandatory	PASS with Comments
12.1 TI iSER	Beta	Not Available
12.2 TI NFS over RDMA	Beta	Not Tested
12.3 TI RDS	Beta	FAIL
12.4 TI SDP	N/A	Not Applicable
12.5 TI uDAPL	Beta	FAIL
12.6-7 TI Basic RDMA Interop	Beta	FAIL
12.8-9 TI RDMA Stress	Beta	Not Available
12.10 TI MPI - Intel	Beta	Not Tested
12.11 TI MPI - Open MPI - homogeneous	Mandatory	PASS
12.11 TI MPI - Open MPI - heterogeneous	Beta	FAIL
12.12 TI MPI - MVAPICH2 - homogeneous	Mandatory	PASS
12.12 TI MPI - MVAPICH2 - heterogeneous	Beta	FAIL

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

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: 4B 9E 65 5C 58 2A 39 80 84 EF 7C 0A BC ED 1E BF

SHA-1 Fingerprint: 02 CB 7B 8F F1 EC 59 21 DE 3F A2 1B 66 06 B8 09 12 D9 DD 0E

Report Revision History

- 1.0 Initial release
- 2.0 Updated with new Adobe signature keys.

Table 1: Result Key

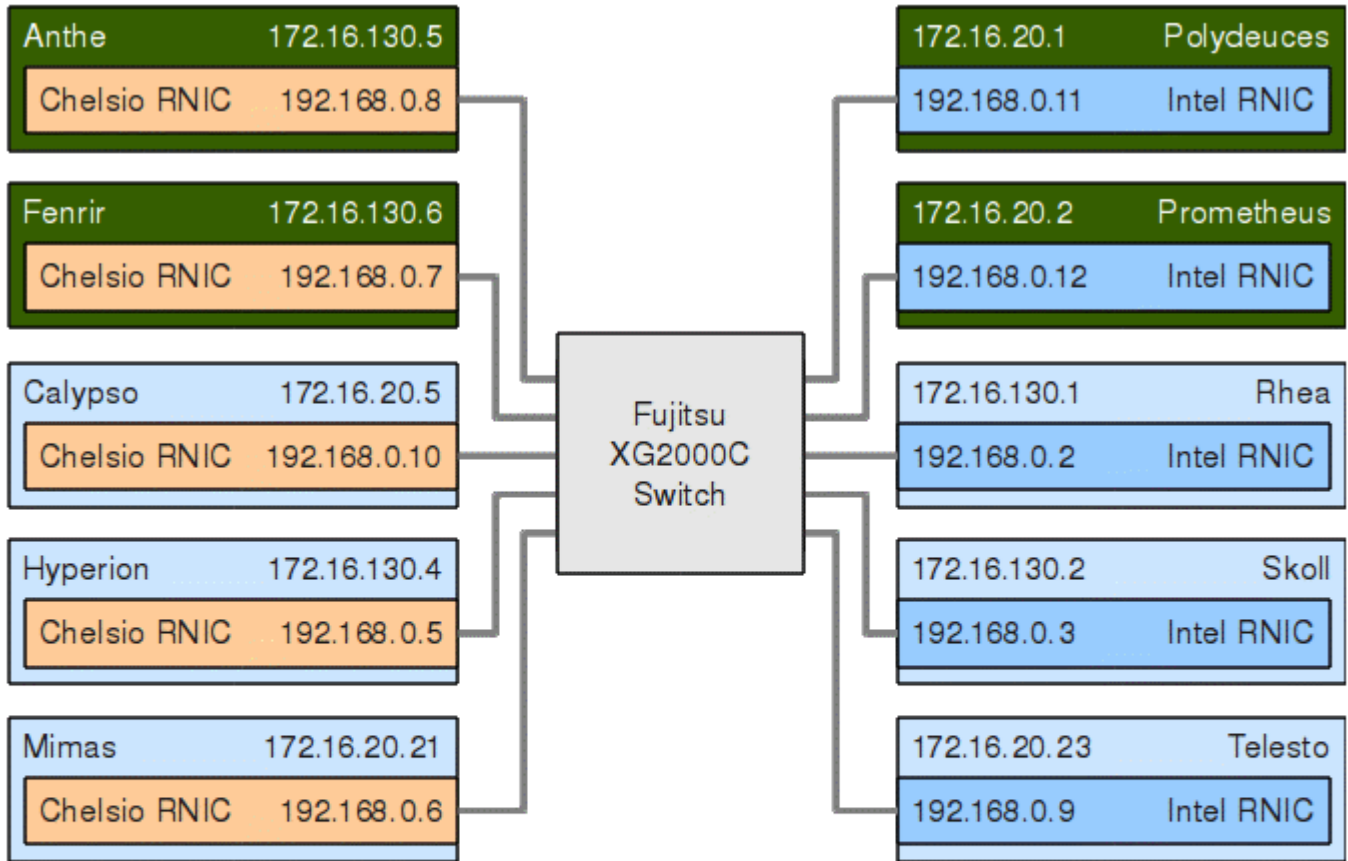
The following table contains possible results and their meanings:

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.
Conditional PASS	The DUT was observed to exhibit conformant behavior under particular conditions.
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 Information

DUT #1 Details			
Manufacturer	Intel (NetEffect)	Firmware Version	3.18
Model	NE0200101E3CX0104	Hardware Version	E59737-001 RV01
Driver	iw_nes version 1.5.0.0	Located in Hosts	Polydeuces (8 cores), Prometheus (8 cores), Rhea (4 cores), Skoll (4 cores), Telesto (4 cores)
Additional Comments/Notes			
Unless otherwise specified, both the Chelsio parameter file /sys/module/iw_cxgb3/parameters/peer2peer on Chelsio hosts and the Intel parameter file /sys/module/iw_nes/parameters/send_first on Intel hosts contain "0".			

Test Setup:



Host hardware key:

AMD Opteron
8 cores, 2.1 GHz
8 GB RAM

Intel Core 2 (Supermicro)
4 cores, 2.7 GHz
4 GB RAM

Device Test Summary Results:

Test Number and Name	Parts	Summary	Result
11.1 Ethernet Link Initialization		error message	PASS with Comments
Details			
rping successfully transmits data across the RDMA channel. However, an error message appears. Executing rping from Chelsio as client to Intel as server displays "cq completion failed status 5" on the Intel/server side.			

Test Number and Name	Parts	Summary	Result
11.2 Ethernet Fabric Initialization		No switches to test	Not Available
Details			

Test Number and Name	Parts	Summary	Result
11.3 Ethernet Fabric Reconvergence		No switches to test	Not Available
Details			

Test Number and Name	Parts	Summary	Result
11.4 Ethernet Fabric Failover		No switches to test	Not Available
Details			

Test Number and Name	Parts	Summary	Result
11.5 iWARP Connectivity	Groups 1-7	unsupported features	PASS with Comments
Details			
Comments from the separate <i>iWARP Consortium Interoperability Report</i> pertaining to several subtests within the iWARP Connectivity test: <ol style="list-style-type: none"> 1.The OFED software stack does not support SendINV and SendSEINV. 2.CX4 and 10G optics do not support speed negotiation. 3.The OFED software stack does not support the ability to configure markers or CRC at runtime. 			

Test Number and Name	Parts	Summary	Result
12.1 TI iSER		No iSER target	Not Available
Details			
There were no iSER targets in the test cluster.			

Test Number and Name	Parts	Summary	Result
12.2 TI NFS over RDMA		Insufficient time	Not Tested
Details			

Test Number and Name	Parts	Summary	Result
12.3 TI RDS	rds-ping	no output	FAIL
	rds-stress	Outputs all zeros.	FAIL
Details			

Test Number and Name	Parts	Summary	Result
<p>For this test the Chelsio parameter file /sys/module/iw_cxgb3/parameters/peer2peer on the Chelsio host contained "1".</p> <p>rds-ping: Homogeneous passes. Heterogeneous produces no output but ends after the appropriate amount of time (3 seconds for a ping count of 3).</p> <p>rds-stress : Homogeneous and heterogeneous behave as if passing, but output values are all zero.</p>			

Test Number and Name	Parts	Summary	Result
12.4 TI SDP		no SDP support	Not Applicable
Details			
Current licensing restrictions prevent SDP support in iWARP RNIC solutions.			

Test Number and Name	Parts	Summary	Result
12.5 TI uDAPL		Hangs	FAIL
Details			
<p>For this test the Chelsio parameter file /sys/module/iw_cxgb3/parameters/peer2peer on the Chelsio host contained "1".</p> <p>Homogeneous passes. Heterogeneous Chelsio (as client) to Intel (as server) hangs on test 3.2. Heterogeneous Intel (as client) to Chelsio (as server) hangs on test 1.1.</p>			

Test Number and Name	Parts	Summary	Result
12.6-7 TI Basic RDMA Interop		Hangs	FAIL
Details			
<p>Chelsio as server, Intel as client: pass. During the "Small RDMA Write" test, 3 of the 10 runs produced the message pp_send_done: bad wc status 5 on the server side but still passed. During the "Large RDMA Write" test, 2 of the 10 runs produced the same message.</p> <p>Intel as server, Chelsio as client: fail. "Small RDMA Write" passes. "Large RDMA Write" fails (pauses interminably) with the message pp_wait_for_done: bad wc status 5 on the server side.</p>			

Test Number and Name	Parts	Summary	Result
12.8-9 TI RDMA Stress		No switches to test	Not Available
Details			

Test Number and Name	Parts	Summary	Result
12.10 TI MPI - Intel		Insufficient time	Not Tested
Details			

Test Number and Name	Parts	Summary	Result

Test Number and Name	Parts	Summary	Result
12.11 TI MPI - Open MPI - homogeneous			PASS
Details			
MPI tests use all five hosts/RNICs and all 28 cores on those five hosts.			

Test Number and Name	Parts	Summary	Result
12.11 TI MPI - Open MPI - heterogeneous		Receive queues incompatible.	FAIL
Details			
The test aborts with the following message. The Open MPI receive queue configuration for the OpenFabrics devices on two nodes are incompatible...			

Test Number and Name	Parts	Summary	Result
12.12 TI MPI - MVAPICH2 - homogeneous			PASS
Details			
MPI tests use all five hosts/RNICs and all 28 cores on those five hosts.			

Test Number and Name	Parts	Summary	Result
12.12 TI MPI - MVAPICH2 - heterogeneous		Aborts shortly after starting.	FAIL
Details			
The test begins normally, but in the first test (PingPong) after the message size exceeds 32,768 bytes the test aborts with the following message. Got completion with error 5...			