



# 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

Chris Calandro  
 Chelsio Communications  
 370 San Aleso Avenue #100  
 Sunnyvale, CA 94085

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):

*Chelsio Communications S310E-CXA Adapter*

The test suite used for the tests described in this report is *OFA-IWG Interoperability Test Plan Release 1.35* (July 27, 2010). 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](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	<b>PASS with Comments</b>
11.2 Ethernet Fabric Initialization	Mandatory	<b>Not Available</b>
11.5 iWARP Connectivity	Mandatory	<b>PASS with Comments</b>
12.11 Open MPI - homogeneous	Mandatory	<b>PASS</b>
12.12 MVAPICH2 - homogeneous	Mandatory	<b>PASS</b>

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

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](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

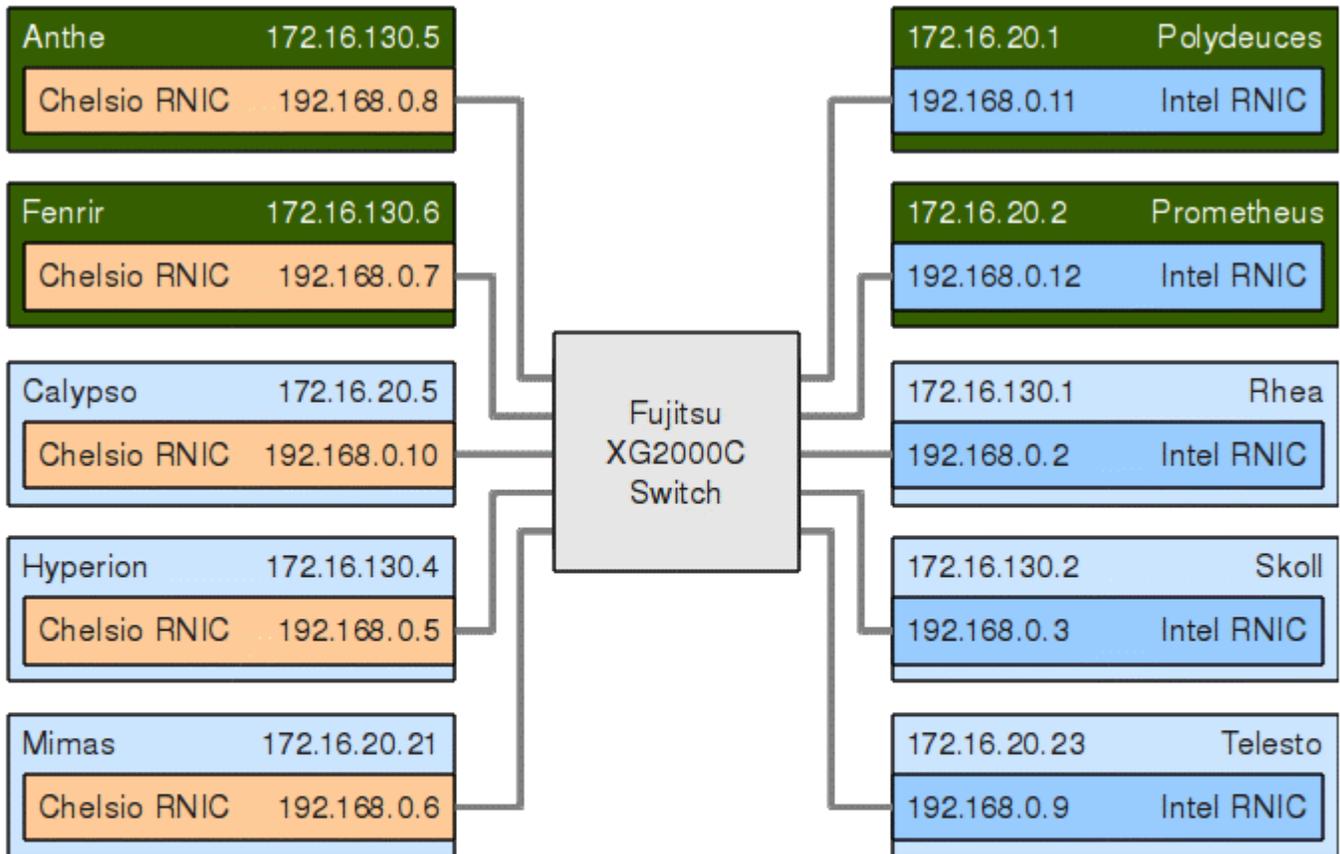
The following table contains possible results and their meanings:

Result	Description
<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>Conditional PASS</b>	The DUT was observed to exhibit conformant behavior under particular conditions.
<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 Information**

<b>DUT #1 Details</b>			
Manufacturer	Chelsio Communications	Firmware Version	T 7.10.0 TP 1.1.0
Model	S310E-CXA	Hardware Version	T3C
Driver	cxgb3 version 1.1.2-ofed	Located in Hosts	Anthe (8 cores), Fenrir (8 cores), Calypso (4 cores), Hyperion (4 cores), Mimas (4 cores)
<b>Additional Comments/Notes</b>			
Unless otherwise specified, 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	<b>PASS with Comments</b>
<b>Details</b>			
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	<b>Not Available</b>
<b>Details</b>			

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

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

Test Number and Name	Parts	Summary	Result
11.5 iWARP Connectivity	Groups 1-7	unsupported features	<b>PASS with Comments</b>
<b>Details</b>			
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"> <li>1.The OFED software stack does not support SendINV and SendSEINV.</li> <li>2.CX4 and 10G optics do not support speed negotiation.</li> <li>3.The OFED software stack does not support the ability to configure markers or CRC at runtime.</li> </ol>			

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

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

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

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 passes. Heterogeneous behaves as if passing, but output values are all zero.</p>			

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

Test Number and Name	Parts	Summary	Result
12.5 TI uDAPL		Hangs	<b>FAIL</b>
<b>Details</b>			
<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	<b>FAIL</b>
<b>Details</b>			
<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	<b>Not Available</b>
<b>Details</b>			

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

Test Number and Name	Parts	Summary	Result

Test Number and Name	Parts	Summary	Result
12.11 TI MPI - Open MPI - homogeneous			<b>PASS</b>
<b>Details</b>			
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.	<b>FAIL</b>
<b>Details</b>			
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			<b>PASS</b>
<b>Details</b>			
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.	<b>FAIL</b>
<b>Details</b>			
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...			