



# OpenFabrics Alliance

## Interoperability Working Group (OFA-IWG)

### June 2009 Logo Validation 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)

Bob Dugan  
Chelsio Communications  
370 San Aleso Avenue #100  
Sunnyvale, CA 94085

August 11, 2009  
Report Revision 1.2  
OFED Version 1.4.1

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 1.28* (March 27, 2009 DRAFT). It is available at the OFA website :

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

Log files are available here:

<ftp://ftp.iol.unh.edu/pub/ofa/>

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
11.5 iWARP Connectivity	Mandatory	<b>PASS with Comments</b>

Testing Completed 7/23/2009

Jeff Laird  
[laird@iol.unh.edu](mailto:laird@iol.unh.edu)

Review Completed 8/11/2009  
Mikkel Hagen

[mhagen@iol.unh.edu](mailto:mhagen@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	Beta	<b>Not Tested</b>
11.2 Ethernet Fabric Initialization	Beta	<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 Operations	Beta	<b>FAIL</b>
12.10 TI MPI - HP	Beta	<b>Not Tested</b>
12.11 TI MPI - Intel	Beta	<b>FAIL</b>
12.12 TI MPI - Open	Beta	<b>FAIL</b>
12.13 TI MPI - OSU	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: F6E2 1B99 28AD 0D25 E77E ADE5 479A 1E05

SHA-1 Fingerprint: AD30 8B08 DD3B B2E3 9362 46E9 3427 BE47 1D49 890B

## Report Revision History

1.0 Initial Release

1.1 Update for peer2peer issue

1.2 Added mandatory-only results table and note about beta tests before all-tests results table. Changed SDP to N/A.

## 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 Information**

DUT #1 Details			
Manufacturer	Chelsio Communications	Firmware Version	T 7.1.0 TP 1.1.0
Model	S310E-CXA	Hardware Version	T3B2
Driver	cxgb3 version 1.0-ofed	Located in Hosts	Polydeuces, Prometheus
Additional Comments/Notes			

**Test Setup:**




## Device Test Summary Results:

Test Number and Name	Parts	Summary	Result
11.1 Ethernet Link Initialization		Insufficient time	<b>Not Tested</b>
<b>Details</b>			

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	Trouble with sends	<b>PASS with Comments</b>
<b>Details</b> 			
<p>When the Chelsio parameter file  /sys/module/iw_cxgb3/parameters/peer2peer  contains "1" on all hosts, tests 2.3 through 2.6 and tests 7.3 through 7.6 (sends) freeze during the ninth send when Intel is sending and Chelsio is receiving. In test 2 that is the ninth send of ten. In test 7 that is the ninth send of 1,000.</p> <p>All tests pass when peer2peer contains "0".</p> <p>Chelsio-to-Chelsio passes regardless of whether peer2peer contains "0" or "1".</p>			

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	Hangs	<b>FAIL</b>
	rds-stress	Hangs	<b>FAIL</b>
<b>Details</b>			
<p>rds-ping: Chelsio-to-Chelsio passes, whether peer2peer is set to 0 or 1. Chelsio-Intel always hangs. When peer2peer=1 and Intel pings Chelsio this hanging does not cause further problems, but when peer2peer=0 or when Chelsio pings Intel this hanging causes Intel-Intel links to stop responding to pings, although they can still transmit pings. The hosts must then be rebooted since the RDS module is not removable.</p> <p>rds-stress : Chelsio-to-Chelsio passes, whether peer2peer is set to 0 or 1. Chelsio- Intel output is like the successful Chelsio-Chelsio output except filled with zeros.  Additionally, after each run of rds-stress, rds-ping is required to free up the port so that rds-stress can run successfully again on the same port. After Intel-Intel rds-stress is run unsuccessfully then rds-ping is run, rds-ping sometimes hangs and sometimes does not.</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>Whereas Chelsio-Chelsio passed in the previous logo event, it fails this time. Strangely, it passes in one direction (Prometheus as client) but not the other (Polydeuces as client), where it hangs on all but the first two tests. Chelsio-Intel with Chelsio as client hangs tests 3.2 through 3.4. Chelsio-Intel with Intel as client gives “length mismatch” error.</p>			

Test Number and Name	Parts	Summary	Result
12.6-7 TI Basic RDMA Interop	bw, lat	Hangs	<b>FAIL</b>
<b>Details</b>			
<p>Chelsio-Chelsio passes the bandwidth test (ib_rdma_bw) but fails the latency test (ib_rdma_lat) with no visible server activity and “Couldn’t create QP” message on the client.  Chelsio-Intel passes the bandwidth test when Chelsio is the server but hangs when Chelsio is the client. Chelsio-Intel with Chelsio as the server fails the latency test with no visible client activity and “Couldn’t create QP” message on the server.  Chelsio-Intel with Chelsio as the client fails the latency test with no visible server activity and “Couldn’t create QP” message on the client.</p> <p>All of the above is with peer2peer=0. With peer2peer=1 the results are the same except Chelsio-Intel with Chelsio as server hangs the bandwidth test.</p>			

Test Number and Name	Parts	Summary	Result
12.8-9 TI RDMA Operations		Same as previous test	<b>FAIL</b>
<b>Details</b>			
There is no difference in the test plan between this test and the previous test.			

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

Test Number and Name	Parts	Summary	Result
12.11 TI MPI - Intel		Timed out	<b>FAIL</b>
<b>Details</b>			
Chelsio-Chelsio timed out before finishing. The test-plan time-out period is 360 seconds. Even with that extended to 720 seconds Chelsio-Chelsio still timed out. Chelsio-Intel timed out before starting.			

Test Number and Name	Parts	Summary	Result
12.12 TI MPI - Open		Some success with no RDMA	<b>FAIL</b>
<b>Details</b>			
With no RDMA (i.e., no “btl openib,self” in the mpirun command line) Chelsio-Chelsio passed and Chelsio-Intel hangs. With RDMA Chelsio-Chelsio and Chelsio-Intel fail with error message, “At least one pair of MPI processes are unable to reach each other...”.			

Test Number and Name	Parts	Summary	Result
12.13 TI MPI - OSU		Hangs	<b>FAIL</b>
<b>Details</b>			
This applies to MVAPICH2 only since MVAPICH does not support iWARP. Chelsio-Chelsio passed, but Chelsio-Intel hangs with no output for over three minutes.			