



OpenFabrics Alliance

Interoperability Logo Group (OFILG)

OFED 3.5-2 Interoperability Report

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

To Whom It May Concern

Date: 25 March 2014
Report Revision: 1.0
OFED Version Tested: 3.5-2
Operating System on Compute Nodes: Scientific Linux 6.4

Enclosed is an interoperability overview of RDMA capable devices tested during the January 2014 Logo Event. Nineteen different devices spanning six device classes and three RDMA transports were tested during this Event. Of the nineteen devices tested during this Logo Event, fifteen were found to pass the mandatory interoperability tests required for the OpenFabrics Interoperability Logo. This document is meant to highlight the performance of member companies' products while alerting the industry to problems that were uncovered with OFED 3.5-2.

The test suite referenced in this report is available at the IOL website. Release 1.49 (2013-Nov-5) was used.

<https://iol.unh.edu/ofatestplan>

The following table outlines the overall status of OFED 3.5-2 interoperability across iWARP, RoCE, and InfiniBand transports as per the Test Plan referenced above.

Test Procedures	iWARP	RoCE	InfiniBand
11.1: Link Initialization	No OFED Issues Discovered	No OFED Issues Discovered	No OFED Issues Discovered
11.2: IB Fabric Initialization	Not Applicable	Not Applicable	No OFED Issues Discovered
11.3: iPoIB Connected Mode	Not Applicable	Not Applicable	No OFED Issues Discovered
11.4: iPoIB Datagram Mode	Not Applicable	Not Applicable	No OFED Issues Discovered
11.5: SM Failover and Handover	Not Applicable	Not Applicable	No OFED Issues Discovered
11.6: SRP	Not Applicable	Not Available	No OFED Issues Discovered
13.1: TI iSER	Not Available	Not Available	Not Available
13.2: TI NFS over RDMA	Not Available	OFED Issue Bug 2449	No OFED Issues Discovered
13.4: TI uDAPL	No OFED Issues Discovered	No OFED Issues Discovered	No OFED Issues Discovered
13.5: TI RDMA Basic Interop	OFED Issue Bug 2457	OFED Issue Bug 2457	OFED Issue Bug 2457
13.6: TI RDMA Stress	No OFED Issues Discovered	Not Tested	No OFED Issues Discovered
13.8: TI MPI – Open MPI	No OFED Issues Discovered	No OFED Issues Discovered	No OFED Issues Discovered

Executive Summary

The goal of the OpenFabrics Interoperability Logo Program is to give assurance to customers, in the form of the OpenFabrics Interoperability Logo, that products on the [Logo List](#) using the OpenFabrics Enterprise Distribution will be interoperable with each other. The grant of the OpenFabrics Interoperability Logo signifies that a product has successfully passed all Mandatory tests during a test event pair and further signifies that the product is interoperable with all other products on the Logo List. By examining the Logo List, customers can be assured that products selected from the List will be compatible with each other, saving deployment time and resources. Each individual product report makes note of any issues that were discovered with the product while executing the test plan.

The Logo Program tests the multi-vendor interoperability of products that employ Remote Direct Memory Access (RDMA) using the OpenFabrics Enterprise Distribution (OFED). The Logo Program is split into two test event pairs:

- The Interop Debug Event
- The Interop GA Event

The Interop Debug Event is held at the UNH-IOL in early April and again in early October. The goal of the Interop Debug Event is to have representatives from OFILG member companies present while testing is being performed so that any and all issues discovered with either the member company's product(s) or OFED at large are reported and resolved in a timely manner. Upon completion of the Interop Debug Event, our members have thirty days to submit bug fixes and firmware updates to enable their products to interoperate with link partners from other companies.

Upon release of the General Availability (GA) release of OFED, the Interop GA Event starts. The goal of the Interop GA Event is to verify multi-vendor interoperability of our members' products using the latest available firmware and the GA release of OFED. The requirements for the OpenFabrics Interoperability Logo are defined in the [Test Plan](#) and [Logo Program](#). Additional general information about the program can be found at our [website](#).

Grant of the OpenFabrics Interoperability Logo signifies that a product has undergone 3-5 months of testing and validation, and that the member companies involved are committed to the multi-vendor interoperability of their respective RDMA technologies and the OpenFabrics Enterprise Distribution. During the January 2014 Interop Logo Event, 8,409 individual command line tests were performed, in addition to the testing done in the months after the Interop Debug Event leading up to the GA Event.

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:

http://www.iol.unh.edu/certifyDoc/certificates_and_fingerprints.php

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: 41 1E 00 9F 79 4D 02 EF E6 95 65 57 A4 71 4F 9F
 SHA-1 Fingerprint: 44 51 9E 22 66 59 1A D3 A1 F9 0B EE BD 01 90 80 BE 61 A4 A8

Report Revision History

- V1.0 Initial working copy, shared with OFILG for comments

Configuration Files

Description	Attachment
OFED 3.5-2 Configuration File	

Result Key

The following table contains possible results and their meanings:

Result:	Description:
No OFED Issues Discovered	The DUTs were observed to exhibit valid behavior.
OFED Issue Bug #####	The DUTs were observed to exhibit invalid behavior due to a discovered and reported bug in OFED.
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.
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.

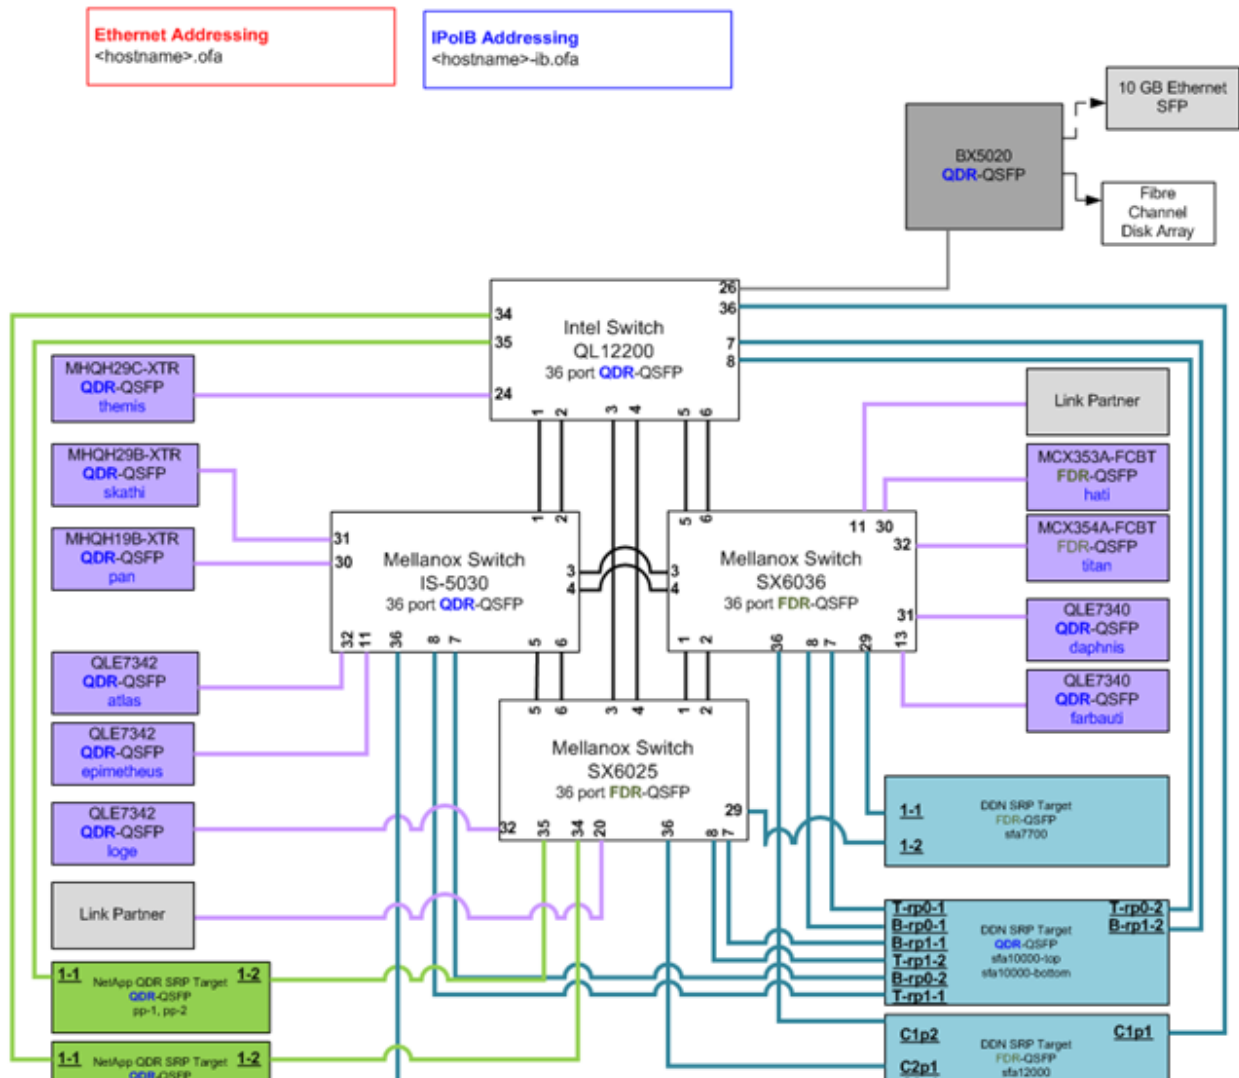


Figure 2: The IW fabric configuration utilized for any tests requiring a multi-switch configuration is shown below.

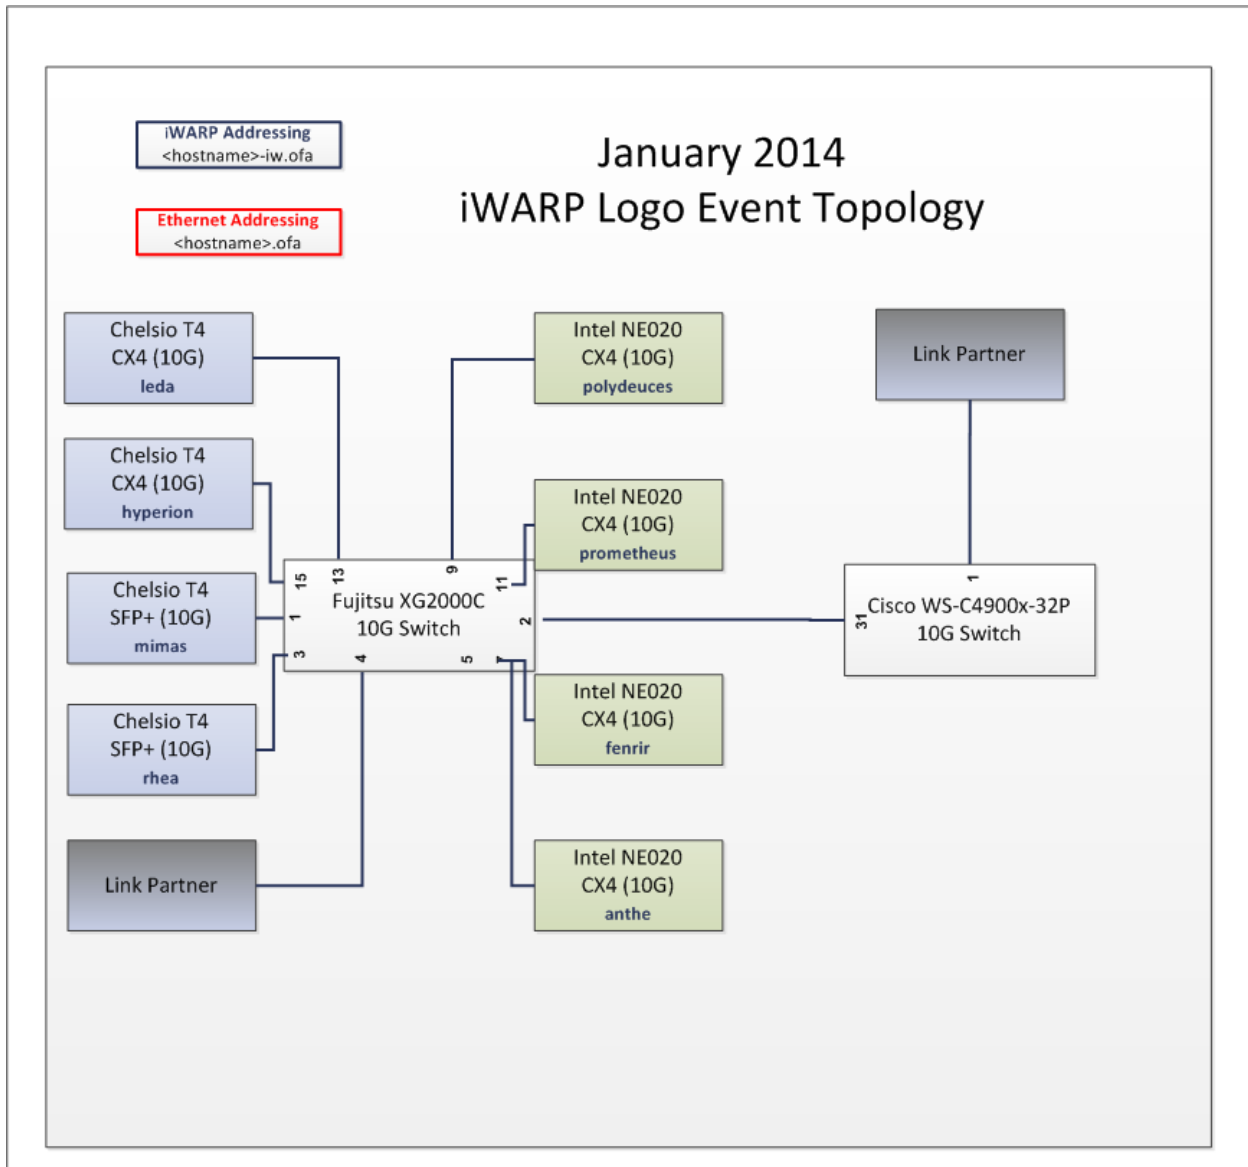
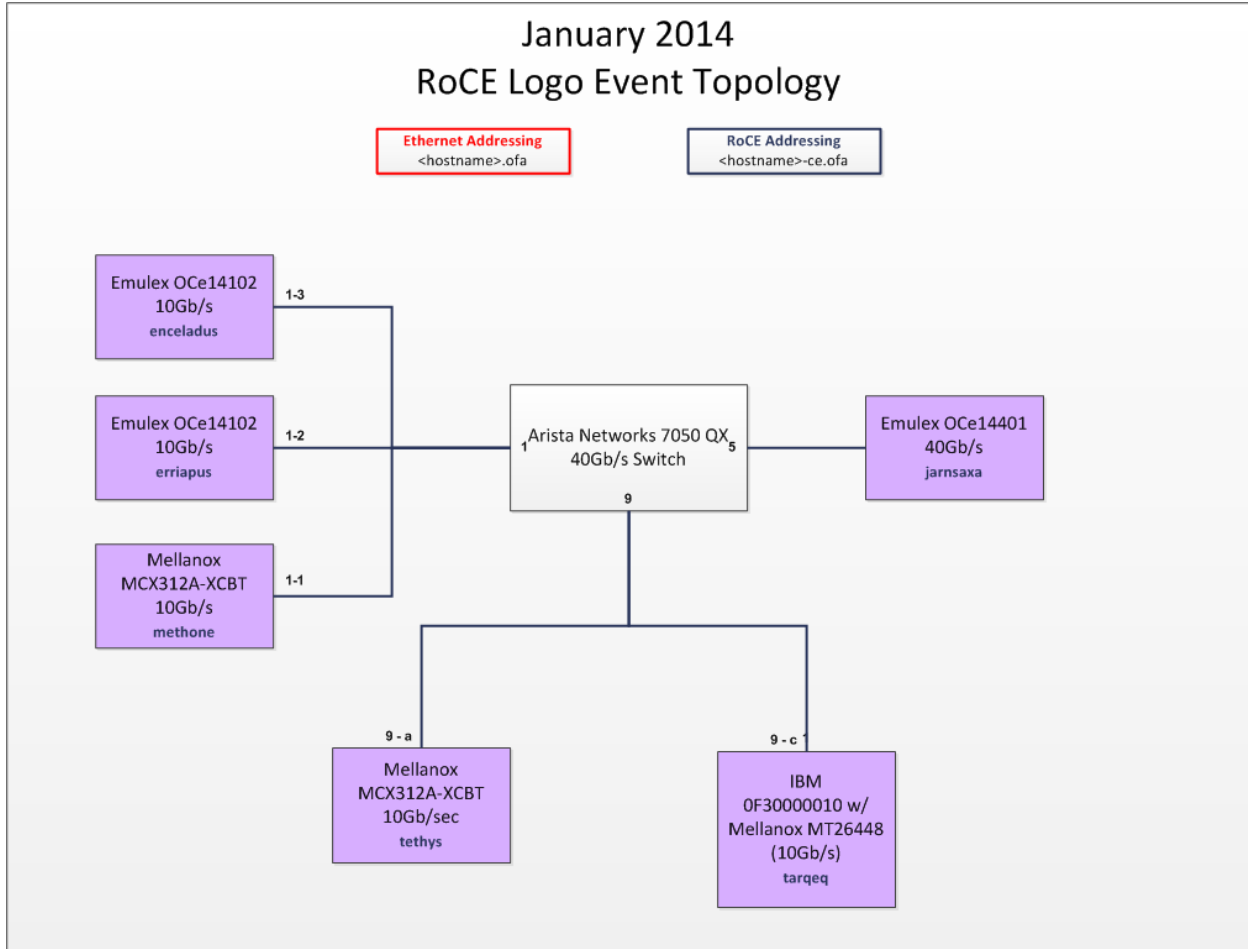


Figure 3: The RoCE fabric configuration utilized for any tests requiring a multi-switch configuration is shown below.



Device Test Results – Transport Independent

The following tests apply to all transports: InfiniBand, iWARP and RoCE.
Please see the [test plan version 1.49](#) for detailed descriptions of each test.

13.1: TI iSER				
Purpose:	To verify that the DUT is capable of reading and writing to and from remote iSCSI block storage using RDMA.			
Transport	<i>InfiniBand</i>	Part A	<i>TI iSERusing OFED</i>	Not Tested
Comments on Test Procedure: None				
Result Discussion: This test was not performed, as there are no devices that support the iSER test procedure present in the event topology.				
Transport	<i>iWARP</i>	Part A	<i>TI iSERusing OFED</i>	Not Tested
Comments on Test Procedure: None				
Result Discussion: This test was not performed, as there are no devices that support the iSER test procedure present in the event topology.				
Transport	<i>RoCE</i>	Part A	<i>TI iSERusing OFED</i>	Not Tested
Comments on Test Procedure: None				
Result Discussion: This test was not performed, as there are no devices that support the iSER test procedure present in the event topology.				

13.2: TI NFS over RDMA				
Purpose:	To verify that the DUT is capable of exporting, mounting, reading, and writing to and from NFS shares using RDMA.			
Transport	<i>InfiniBand</i>	Part A	<i>TI NFS over RDMAusing OFED</i>	No OFED Issues Discovered
Comments on Test Procedure: None				
Result Discussion: All devices that were granted the OFILG Logo during the January 2014 Logo Event were found to properly act as a client and server in an NFSoRDMA client server pair.				
Transport	<i>iWARP</i>	Part A	<i>TI NFS over RDMAusing OFED</i>	Not Available
Comments on Test Procedure: None				
Result Discussion: This test is not currently required for the OFILG Logo due to its Beta status.				
Transport	<i>RoCE</i>	Part A	<i>TI NFS over RDMAusing OFED</i>	OFED issue discovered - BUG 2449
Comments on Test Procedure: None				
Result Discussion: During the January 2014 Interoperability Logo Event, a bug was discovered when interoperating between big and little endian architectures. See OFED Bug 2449 for additional information.				

13.4: TI uDAPL				
Purpose:	To verify that the DUT is capable of utilizing the User-Level Direct Access Transport APIs.			
Transport	<i>InfiniBand</i>	Part A	<i>TI uDAPLTEST Commands using OFED</i>	No OFED Issues Discovered
Comments on Test Procedure: None				
Result Discussion: All devices that were granted the OFILG Logo during the January 2014 Logo Event were found to properly utilize the User-Level Direct Access Transport APIs.				
Transport	<i>iWARP</i>	Part A	<i>TI uDAPLTEST Commands using OFED</i>	No OFED Issues Discovered
Comments on Test Procedure: None				
Result Discussion: All devices that were granted the OFILG Logo during the January 2014 Logo Event were found to properly utilize the User-Level Direct Access Transport APIs.				
Transport	<i>RoCE</i>	Part A	<i>TI uDAPLTEST Commands using OFED</i>	No OFED Issues Discovered
Comments on Test Procedure: None				
Result Discussion: All devices that were granted the OFILG Logo during the January 2014 Logo Event were found to properly utilize the User-Level Direct Access Transport APIs.				

13.5: TI RDMA Basic Interoperability				
Purpose:	To demonstrate the ability of endpoints to exchange core RDMA operations across a simple network path.			
Transport	<i>InfiniBand</i>	Part A	<i>TI RDMA Basic Interop</i>	OFED Issue Discovered - BUG 2457
Comments on Test Procedure: None				
Result Discussion: During the January 2014 Interoperability Logo Event, a bug was discovered when executing RDMA sends of message sizes less than 240 bytes. See OFED Bug 2457 for additional information. This issue was resolved in OFED 3.12				
Transport	<i>iWARP</i>	Part A	<i>TI RDMA Basic Interop</i>	OFED Issue Discovered - BUG 2457
Comments on Test Procedure: None				
Result Discussion: During the January 2014 Interoperability Logo Event, a bug was discovered when executing RDMA sends of message sizes less than 240 bytes. See OFED Bug 2457 for additional information. This issue was resolved in OFED 3.12				
Transport	<i>RoCE</i>	Part A	<i>TI RDMA Basic Interop</i>	OFED Issue Discovered - BUG 2457
Comments on Test Procedure: None				
Result Discussion: During the January 2014 Interoperability Logo Event, a bug was discovered when executing RDMA sends of message sizes less than 240 bytes. See OFED Bug 2457 for additional information. This issue was resolved in OFED 3.12				

13.6: TI RDMA Stress				
Purpose:	To identify problems that may arise when RDMA operations are performed over interconnected devices in the fabric.			
Transport	<i>InfiniBand</i>	Part A	<i>TI RDMA Stress Test</i>	No OFED Issues Discovered
Comments on Test Procedure: None				
Result Discussion: All devices that were granted the OFILG Logo during the January 2014 Logo Event were found to properly execute RDMA read, write, and send operations.				
Transport	<i>iWARP</i>	Part A	<i>TI RDMA Stress Test</i>	No OFED Issues Discovered
Comments on Test Procedure: None				
Result Discussion: All devices that were granted the OFILG Logo during the January 2014 Logo Event were found to properly execute RDMA read, write, and send operations.				
Transport	<i>RoCE</i>	Part A	<i>TI RDMA Stress Test</i>	Not Tested
Comments on Test Procedure: None				
Result Discussion: This test was not executed during the January 2014 Logo Event due to its Beta status and RoCE fabric constraints. This test will be performed in the future, once additional Ethernet switches are introduced into the topology.				

13.8: TI MPI – OpenMPI				
Purpose:	To verify that endpoints are able to correctly pass messages using an RDMA network, by means of the Intel MPI Benchmarks.			
Transport	<i>InfiniBand</i>	Part A	<i>TI MPI – Open MPI using OFED</i>	No OFED Issues Discovered
Comments on Test Procedure: None				
Result Discussion: All devices that were granted the OFILG Logo during the January 2014 Logo Event were found to properly utilize the Intel Message Passing Interface.				
Transport	<i>iWARP</i>	Part A	<i>TI MPI – Open MPI using OFED</i>	No OFED Issues Discovered
Comments on Test Procedure: None				
Result Discussion: All devices that were granted the OFILG Logo during the January 2014 Logo Event were found to properly utilize the Intel Message Passing Interface.				
Transport	<i>RoCE</i>	Part A	<i>TI MPI – Open MPI using OFED</i>	No OFED Issues Discovered
Comments on Test Procedure: None				
Result Discussion: All devices that were granted the OFILG Logo during the January 2014 Logo Event were found to properly utilize the Intel Message Passing Interface.				

Device Test Results – InfiniBand only

11.1: InfiniBand Link Initialization		
Purpose:	To verify that the DUT is capable of properly initializing a link at the proper width and speed with all other link partners in the fabric.	
Part A	<i>IB Link Initialization using OFED</i>	No OFED Issues Discovered
Comments on Test Procedure: None		
Result Discussion: All devices that were granted the OFILG Logo during the January 2014 Logo Event were found to establish a link with all other Link Partners at the proper speed and width.		

11.2: InfiniBand Fabric Initialization		
Purpose:	To verify that the DUT is capable of properly initializing an InfiniBand fabric using OpenSM.	
Part A	<i>IB Fabric Initialization using OFED</i>	No OFED Issues Discovered
Comments on Test Procedure: None		
Result Discussion: All devices that were granted the OFILG Logo during the January 2014 Logo Event were found properly initialize the InfiniBand fabric when using OpenSM.		

11.3: IPoIB Connected Mode		
Purpose:	To verify that the DUT is capable of properly performing Internet Protocol operations over the InfiniBand transport in Connected mode.	
Part A	<i>IPoIB Connected Mode using OFED</i>	No OFED Issues Discovered
Comments on Test Procedure: None		
Result Discussion: All devices that were granted the OFILG Logo during the January 2014 Logo Event were found properly complete IPoIB operations.		

11.4: IPoIB Datagram Mode		
Purpose:	To verify that the DUT is capable of properly performing Internet Protocol operations over the InfiniBand transport in Datagram mode.	
Part A	<i>IPoIB Datagram Mode using OFED</i>	No OFED Issues Discovered
Comments on Test Procedure: None		
Result Discussion: All devices that were granted the OFILG Logo during the January 2014 Logo Event were found properly complete IPoIB operations.		

11.5: InfiniBand SM Failover and Handover		
Purpose:	To verify that the DUT is capable of properly handling SM priority and state rules.	
Part A	<i>SM Failover and Handover using OFED</i>	No OFED Issues Discovered
Comments on Test Procedure: None		
Result Discussion: All devices that were granted the OFILG Logo during the January 2014 Logo Event were found to properly handle Subnet Manager priority and state rules using the OpenSM.		

11.6: InfiniBand SRP using OFED		
Purpose:	To verify that the DUT is capable of reading and writing to remote block storage devices using RDMA.	
Part A	<i>IB SRP Using OFED Core</i>	No OFED Issues Discovered
Part B	<i>IB SRP Using OFED Extended</i>	No OFED Issues Discovered
Comments on Test Procedure: None		
Result Discussion: All devices that were granted the OFILG Logo during the January 2014 Logo Event were found to properly perform SRP operations.		

Device Test Results – iWARP only

12.1: iWARP Link Initialization		
Purpose:	To verify that the DUT is capable of linking at the correct speed and passing traffic to a Link Partner under nominal (unstressed) conditions.	
Part A	<i>iWARP Link Initialize using OFED</i>	No OFED Issues Discovered
Comments on Test Procedure: None		
Result Discussion: All devices that were granted the OFILG Logo during the January 2014 Logo Event were found to establish a link with all other Link Partners at the proper speed.		

Device Test Results – RoCE only

12.2: RoCE Link Initialization		
Purpose:	To verify that the DUT is capable of linking at the correct speed and passing traffic to a Link Partner under nominal (unstressed) conditions.	
Part A	<i>RoCE Link Initialize using OFED</i>	No OFED Issues Discovered
Comments on Test Procedure: None		
Result Discussion: All devices that were granted the OFILG Logo during the January 2014 Logo Event were found to establish a link with all other Link Partners at the proper speed.		