

OFA-IWG Interop Event March 2009



OPENFABRICS
ALLIANCE

Rupert Dance & Arkady Kanevsky
Sonoma Workshop Presentation

OFA-IWG Charter and Mission



- OpenFabrics Alliance Interoperability Working Group (OFA-IWG)
 - The OFA-IWG is responsible for defining and maintaining the OFA Logo Program.
- Define and develop a comprehensive suite of tests for evaluating product interoperability within the OpenFabrics environment.
- Create new tests when features are added to OFED
- Validate OFED Release Candidates before the GA release
- Host interoperability events in conjunction with the University of New Hampshire Interoperability Lab ([UNH-IOL](#)).
 - Validate the interoperability of products using the OpenFabrics software stack.
 - The [OFA Interoperability Logo](#) is granted to those products which pass all of the mandatory tests.

Event Overview

- OFA IWG Interop Events - [Upcoming](#)
 - March 16th → 20th 2009 – OFED 1.4.1 RC Validation
 - April 23rd → May 1st 2009 – Interop Debug Event
 - May 25th → May 29th 2009 – Logo Validation – OFED 1.4.1 GA
- This will be the 5th OFA-IWG Interop event
- Anticipated participation
 - 10 vendors and 37 devices
- Compliance Testing
 - April 20th → 24th – IBTA Plugfest
 - April 20th → 24th – LNI iWARP Plugfest
- Hosted by University of NH Interoperability Lab (UNH-IOL)

Program Details

- OFA Interoperability Logo Program
 - Enable Customers and users to choose products and versions that are interoperable
 - Enable vendors to demonstrate interoperability by using the Logo on their products and marketing.
 - Details available at <http://www.iol.unh.edu/services/testing/ofa/logo.pdf>
 - Current List of Logo Grants from December 2008
 - <http://www.iol.unh.edu/services/testing/ofa/interoplist/interoplist.php>

Vendor Participation - IB

- InfiniBand Vendors
 - Data Direct Networks, Flextronics, LSI Logic, Mellanox, Obsidian Research, QLogic, Voltaire
- InfiniBand Devices – using OFED 1.4 GA
 - 16 HCAs
 - 5 Switches
 - 2 InfiniBand Range Extenders
 - 3 SRP Targets
 - 1 FC/Ethernet Gateway
 - 3 SMs
 - OFED OpenSM
 - QLogic Switch Based SM
 - Voltaire Switch Based SM

Vendor Participation - iWARP



- iWARP Vendors
 - Chelsio and Intel (NetEffect)
- iWARP Devices - using OFED 1.4 GA
 - Chelsio
 - 3 RNICs – 2 CX4 and one SR
 - Intel (NetEffect)
 - 2 RNICs
 - Fujitsu
 - 10 GbE Switches

Test Setup & Environment

- Host Systems
 - 17 servers donated by Intel
 - 6 servers donated by HP
- Cables Donated to OFA Cluster
 - Amphenol, C&M, Intel, Mellanox, Meritec, Molex, Panduit, W.L. Gore, Zarlink
- RHEL 5.2 (CentOS 5.2)
 - 64 bit x86
 - 64 bit Itanium
- OFED 1.4.1 GA

Test Coverage

➤ InfiniBand

- Link Init & Fabric Init
- IPoIB CM and UD
- iSER & SRP
- Intel MPI, HP MPI, Open MPI & MVAPICH 1 & 2
- RDS & NFS/RDMA
- SDP
- RDMA Interop
- uDAPL

➤ iWARP

- iWARP Interop
- Intel MPI, HP MPI, Open MPI & MVAPICH 1 & 2
- RDMA Interop
- RDS & NFS/RDMA
- SDP
- uDAPL
- Ethernet switch testing (execution dependent on participation of switch vendors)

Test Updates for April 2009

- Updates to IB tests – new automation scripts
 - Link Init, Fabric Init, IPoIB CM, IPoIB UD, SDP, SRP
- MPI tests updated for homogenous testing only
 - All MPIs tests run in homogeneous host and HCA environment
 - HP MPI test updated to the latest HP MPI release
 - Open MPI updated to use version 1.3.1
- TI - RDS - new
 - Testing for Reliable Datagram Sockets
 - Updated to support both IB and iWARP
 - Now uses *rds-ping* and *rds-stress*
- TI - NFS RDMA - new
 - New test procedure for April 2009
 - Uses the Connectathon Test Suite

IB Tests - April 2009



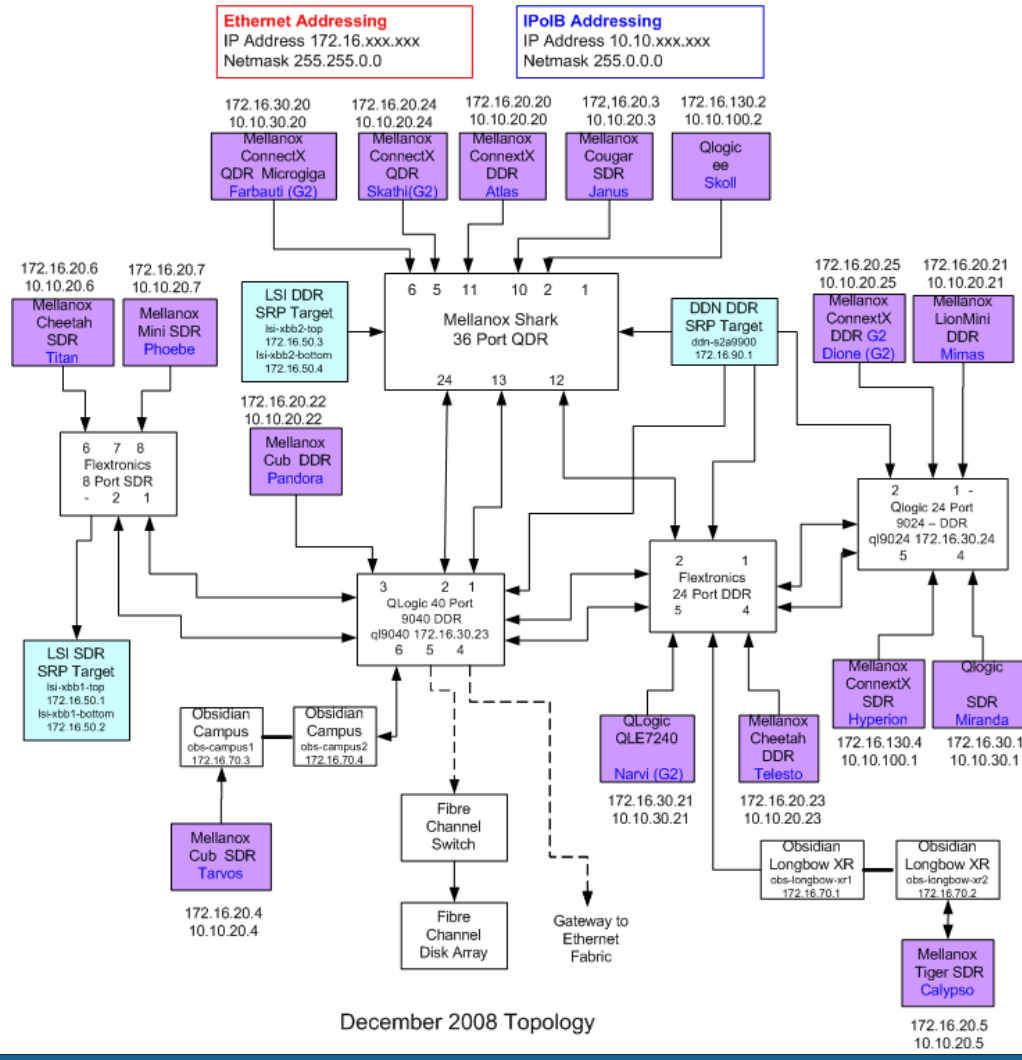
InfiniBand Transport - Test Status for April 2009 Interop Event

Test Procedures	Linux	WinOF
IB Link Initialize	Mandatory	Beta
IB Fabric Initialization	Mandatory	Beta
IB IPoIB Connected Mode	Mandatory	N/A
IB IPoIB Datagram Mode	Mandatory	Beta
IB SM Failover/Handover	Beta	Beta
IB SRP	Mandatory	Beta
IB Ethernet Gateway	Beta	N/A
IB Fibre Channel Gateway	Beta	N/A
Transport Independent	Linux	WinOF
TI iSER	Mandatory	N/A
TI NFS over RDMA	Beta	N/A
TI RDS	Beta	N/A
TI SDP	Mandatory	N/A
TI uDAPL	Beta	Beta
TI MPI - HP	Beta	N/A
TI MPI - Intel	Beta	Beta
TI MPI - Open MPI	Beta	N/A
TI MPI - OSU - MVAPICH 1 & 2	Beta	N/A
TI RDMA Interop	Beta	N/A
TI RDMA Operations	Beta	N/A

Ethernet Tests - April 2009

Ethernet Transport - Test Status for April 2009 Interop Event	
Test Procedures	Linux
Ethernet Link Initialize	Beta
Ethernet Fabric Initialize	Beta
Ethernet Fabric Failover	Beta
Ethernet Fabric Reconvergence	Beta
iWARP Connectivity	Mandatory
Transport Independent	Linux
TI iSER	Mandatory
TI MPI - HP	Beta
TI MPI - Intel	Beta
TI MPI - Open MPI	Beta
TI MPI - OSU - MVAPICH 1 & 2	Beta
TI NFS over RDMA	Beta
TI RDMA Interop	Beta
TI RDMA Operations	Beta
TI RDS	Beta
TI SDP	Mandatory
TI uDAPL	Beta

IB Cluster Topology



WinOF IB Tests - April 2009



- Beta - InfiniBand
 - Link Init & Fabric Init
 - OpenSM
 - IPoIB UD
 - SRP
 - Intel MPI
 - RDMA Interop
 - uDAPL
- Microsoft Engagement

Problems Noted – October 2008



- HCA to HCA link issue
 - Does not occur when HCAs are connected through switches
 - Will be resolved for April 2009 Interop Event
- IPoIB
 - Drops the first packet when size ≥ 8192
 - OFED issue noted in Bug 1287 – open issue to be resolved in OFED 1.4.1
- MPI – unsupported heterogeneous environment
 - MPI fails when running mixed HCAs
 - MPI fails when running more than 38 processes on non-homogenous nodes
- SRP
 - Polarity issues and interconnect issues between HCA and SRP TCA
 - Resolved during Interop Debug event and device passed GA Validation

Notes & Reminders

- Cluster Availability
 - VPN access is granted to all OFILG Members.
 - Week long testing slots are available.
- Next Event Dates - 2009
 - OFED 1.4.1 RC Validation: March 16th – 20th
 - OFA Interop Debug Event: April 23th – May 1st
 - OFA Logo Validation: May 25th – May 29th
 - IBTA Plugfest : April 20th – 24th
 - iWARP Compliance: April 20th – 24th

UNH-IOL OFA – Moving Forward



- Logo Program has achieved the technical goals of the OFA
- Logo Program needs more marketing & promotion
- Original scope of the contract
 - Two Interop Events per year
- Current scope
 - Two RC validation test periods
 - Two Interop Debug events based on RC
 - Two Logo Validation events based on GA
- Number of tests has nearly doubled
- Additional funding needed
 - Increased fees per member
 - Increased membership

Open Invitation for Interop #5



- Distros
 - Novell and Red Hat
- Ethernet Switch Vendors
 - Arista Networks, Force 10, Fujitsu, Fulcrum, Woven
- IB Vendors
 - Gennum, HP, IBM, SGI, Sun, Xsigo, XLoom
- iWARP Vendors
 - NetXen, IBM
- System Vendors
 - Appro, Dell, HP, SuperMicro, Verrari

IBTA Plugfest - Fall 2008



- IBTA IL Compliance Testing September 22 - 26, 2008
 - 18 vendors
 - 16 devices
 - 10 HCAs, 3 switches, 2 Range Extenders, 1 SRP Target
 - 114 cables tested – versus 83 at the last event
 - This was the 14th IBTA Plugfest
- IBTA integrators List
 - This list includes all products that have successfully completed the rigorous testing procedures at IBTA sponsored Plugfests.
 - Current List of devices is available at <http://www.infinibandta.org/itinfo/IL>
- Hosted by Lamprey Networks and University of NH Interoperability Lab (UNH-IOL)
- Next event scheduled for April 20th 2009

IBTA Vendor Participation



- IBTA Plugfest #14 - September 22 - 26, 2008
- InfiniBand Devices
 - Bay Microsystems, Flextronics, LSI Logic, Mellanox, Obsidian Research, QLogic
- Cable Vendors
 - Amphenol, C&M, Cinch, Fujikura, Emcore, Luxtera, Merteck, Molex, Panduit, Tyco, W.L. Gore, Zarlink

IBTA IL Logo Program



- The official logo is the basic identity for the InfiniBand Trade Association certification program.
- The official certified logos can be used on websites, advertising, marketing, and collateral material for those products that have been officially IBTA certified.
- If one wants to use it on collateral material, they must also use it on the product itself or on product packaging.
- All products certified on the September 2008 IL List are eligible.
- Products currently being manufactured with the Logo

IBTA Integrators List Logos

Device Logo



Cable Logos 4X

Cable Logos 12X



Cinch Connectors – IBTA Logo

