



OpenFabrics Alliance Interoperability Program

Authors: Rupert Dance, Nick Wood

Date: 3/26/2012

Agenda

- OFA Interoperability Program Goals & Benefits
- OFA Clusters and Equipment at UNH-IOL
- December 2012 Logo Validation Event
 - Test coverage
 - IB and iWARP Topology
- Logo Validation Results
 - Explanation of Out-of-the-Box testing
- Problems noted during the event
- OFILP Program Summary
- OFILP plans for the future
- OFA member call to action

OFA Interoperability Working Group Charter



- The OFA-IWG defines & maintains the OFA Interoperability Logo Program (OFILP)
- Develop a suite of tests for evaluating interoperability between OFA Software and RDMA products using a variety of operating systems.
- Create new tests when features are added to OFA Software.
- Validate OFA Software during EWG development process and Distro 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.

Interoperability Logo Program - Benefits



- Validate OFED & Distro releases in a diversified & heterogeneous cluster.
- Enable vendors to test new hardware, firmware and software using the latest products available from all vendors.
 - VPN access to the cluster 24*7
- Rapid redeployment of cluster topology (< 30 minutes)
 - CentOS, RedHat, Scientific Linux, SUSE, W2K8 HPC
 - OFED 1.4.x.x, 1.5.x.x, 3.2 and WinOFED 3.0
- Enable **end users** to access the [Logo List](#) that identifies RDMA HW & SW and Distros that have demonstrated interoperability.
- Enable **vendors** to demonstrate interoperability by using the OFA Tested Logo on their products and marketing materials.



New additions in 2011 to OFILP


- Added Testing for RHEL Distributions
- Added Testing for SUSE Distributions
- Added Testing of Scientific Linux
- Added OFED version interoperability
- Added RDMA Interoperability between Distros
- Added Testing for Windows W2K8 & WinOFED

OFA Clusters at the UNH-IOL



- Clusters
 - Linux
 - iWARP
 - InfiniBand
 - Windows
 - InfiniBand
- Server Donations
 - 12 AMD Servers
 - 19 Intel Servers
 - 6 HP Servers
 - 6 OFA Servers
- Resources
 - 32,000 Square Feet
 - 100+ staff & students

OFILG Equipment at the UNH-IOL



- iWARP Hardware Available
 - 4 Chelsio RNICs (CX4 & SFP+)
 - 1 Fujitsu 10 GbE Switch
 - 4 Intel RNICs
- InfiniBand Hardware Available
 - 2 DDN SRP targets
 - 1 DDR, 1 QDR
 - 6 Intel (QLogic) QDR HCAs
 - 2 Intel (QLogic) QDR IB Switches
 - 18 Mellanox HCAs
 - 9 DDR, 7 QDR, 2 FDR
 - 4 Mellanox Switches
 - 2 QDR, 2 FDR
 - 1 Mellanox Ethernet Gateway
 - 2 NetApp SRP Targets
 - 1 DDR, 1 QDR
 - 2 Range Extenders

OFA Interop Event Overview



- Most Recent Event
 - December 2011 -- GA Logo Validation Event - OFED 1.5.4
 - [OpenFabrics Alliance Logo List](#)
- Vendor participation
 - AMD, Chelsio, DataDirect Networks, Intel, Mellanox, Microsoft, NetApp, QLogic, RedHat, SUSE
- Upcoming OFA IWG Interop Events
 - April 23 → April 27 2012 – [Interop Debug Event](#)
 - May 2012 – Logo GA Validation Event - OFED 1.5.4.1
 - This will be the 12th OFA-IWG Interop event
- Hosted by University of NH Interoperability Lab ([UNH-IOL](#))

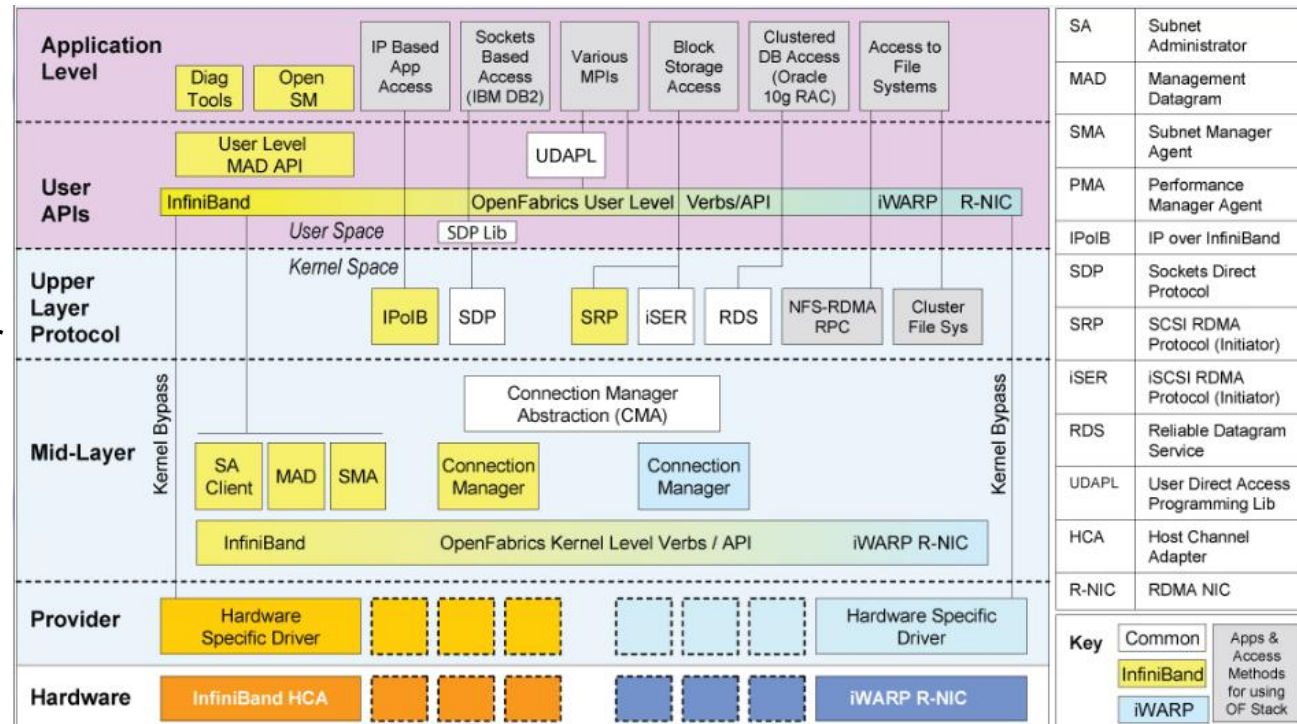
OFA Software Test Coverage

InfiniBand

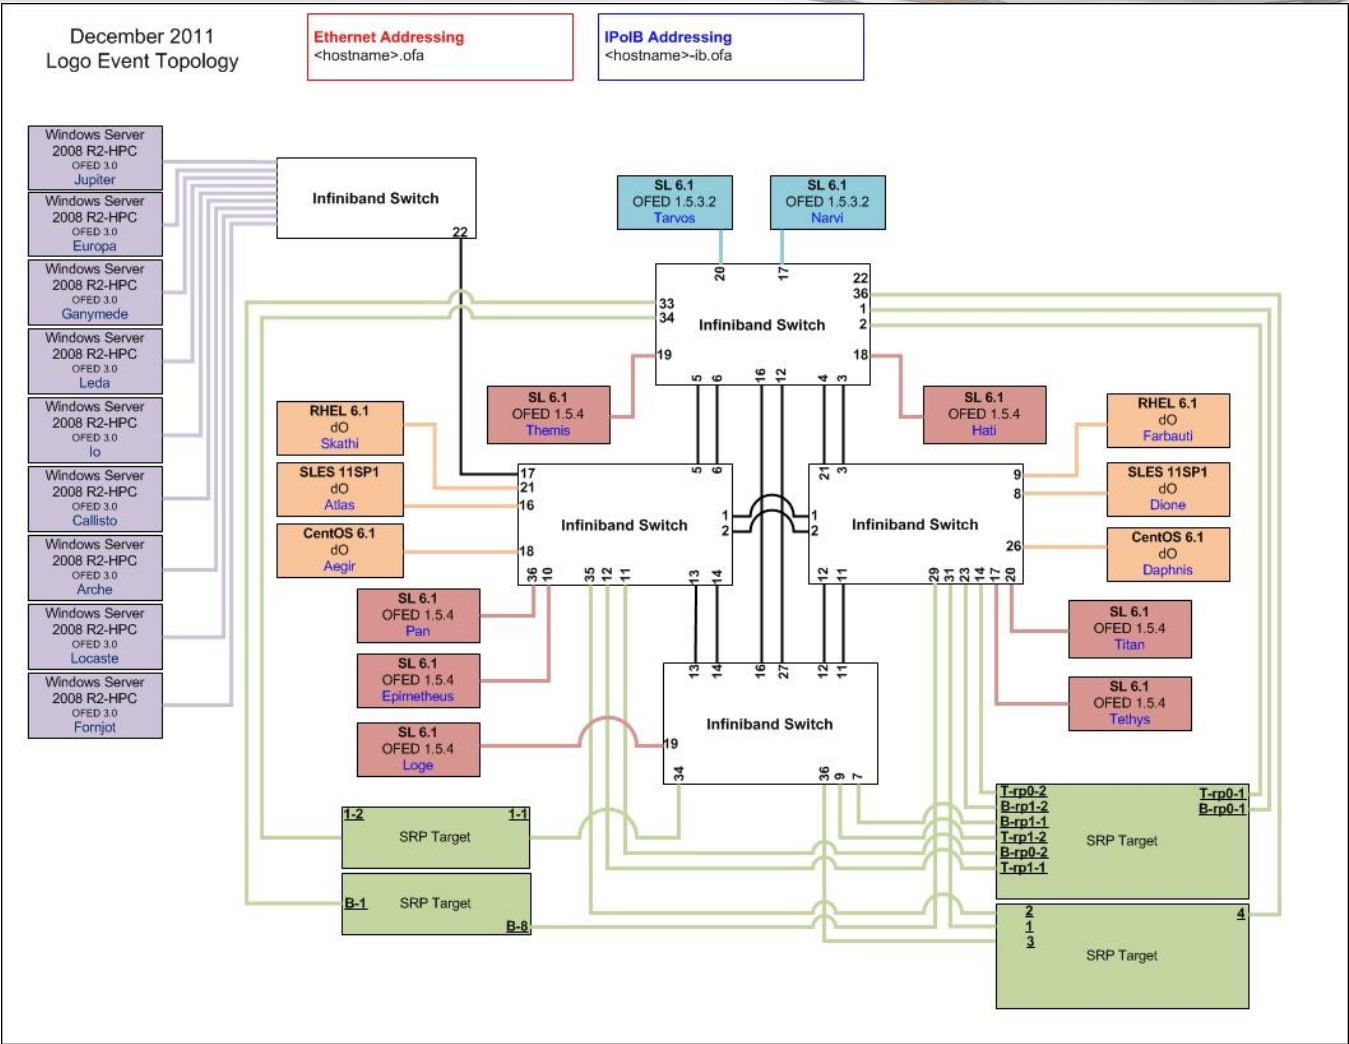
- Link Init & Fabric Init
- IPoIB CM and UD
- Open MPI
- RDMA Interop
- RDS
- NFS/RDMA
- SDP
- SM Failover/Handover
- SRP
- uDAPL

iWARP

- Link Init & Fabric Init
- iWARP Connectivity
- Open MPI
- RDMA Interop
- uDAPL

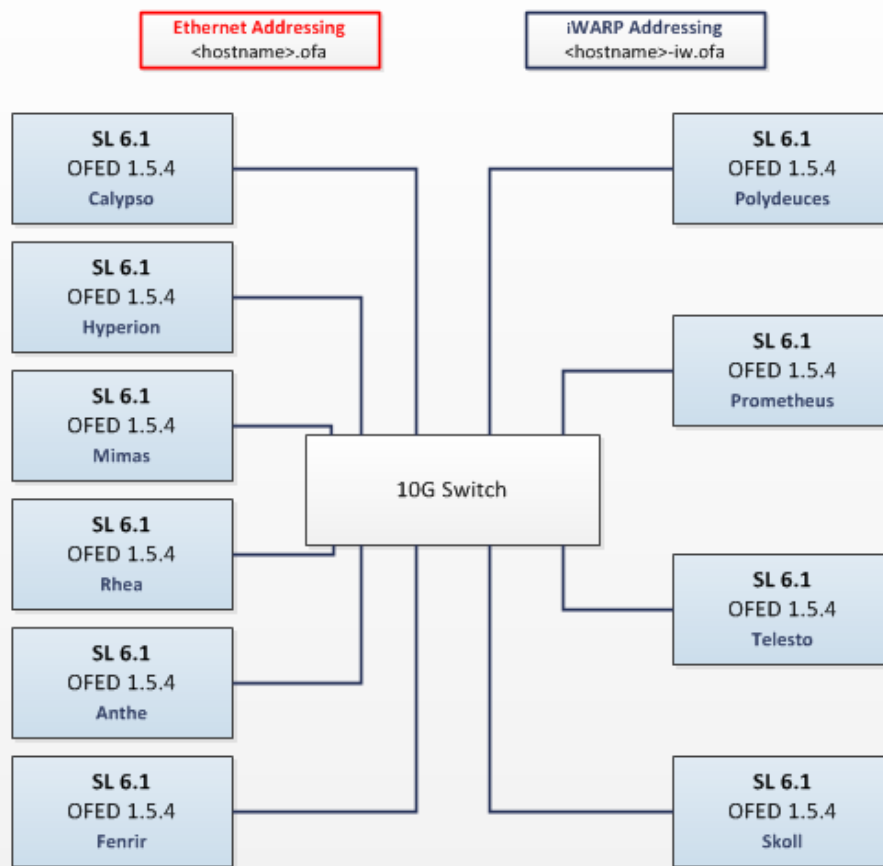


IB Cluster Topology



iWARP Cluster Topology

December 2011 iWARP Logo Event Topology



Out-of-the Box versus OFED

- OFA EWG is now pushing all development work into the kernel
- OFED contains all RPMs needed to deploy complete functionality
- OFA wants to deliver via the Distros
- Need to ensure end user experience out of the box is equivalent to OFED

Logo GA Validation Results

InfiniBand Test procedures	IB-SL 6.1 OFED 1.5.4	IB-SL 6.1 OFED 1.5.3.2	IB-CentOS 6.1 Out-of-Box	IB-RHEL 6.1 Out-of-Box	IB-SLES 11 SP1	W2K8-R2 WinOFED 3.0
IB Link Initialization	FAIL	N/T	N/T	N/T	N/T	N/T
IB Fabric Initialization	PASS	PASS	PASS	PASS	PASS	PASS
IPoIB Connected Mode	PASS	PASS	FAIL	FAIL	FAIL	N/A
IPoIB Datagram Mode	PASS	PASS	PASS	PASS	PASS	N/T
SM Failover and Handover	PASS	PASS	PASS	PASS	PASS	PASS
SCSI Remote Protocol (SRP)	PASS	PASS	PASS	PASS	PASS	N/T
NFS over RDMA	PASS	FAIL	PASS	PASS	FAIL	N/A
RDS	PASS	PASS	FAIL	FAIL	FAIL	N/A
SDP	PASS	PASS	FAIL	FAIL	FAIL	N/A
uDAPL	Pass	PASS	PASS	FAIL	FAIL	N/T
RDMA Basic Interoperability	PASS	PASS	FAIL	FAIL	FAIL	N/A
RDMA Stress	PASS	PASS	FAIL	FAIL	FAIL	N/A
Open MPI	PASS	PASS	FAIL	FAIL	FAIL	N/A
MVAPICH	PASS	PASS	FAIL	FAIL	FAIL	N/A
iWARP Test procedures	IW-SL 6.1 OFED 1.5.4	<p>This is a high level summary and the details are available in the Master Report available for review on request.</p>				
Ethernet Link Initialize	PASS					
Ethernet Fabric Initialize	PASS					
iWARP Connectivity	PASS					
uDAPL	PASS					
RDMA Basic Interoperability	PASS					
Open MPI	PASS					
MVAPICH	PASS					
Legend	Description	Comment				
N/A	Not Applicable					
N/T	Not Tested					
IB-SL 6.1 OFED 1.5.4	InfiniBand using Scientific Linux 6.1 - OFED 1.5.4					
IB-SL 6.1 OFED 1.5.3.2	InfiniBand using Scientific Linux 6.1 - OFED 1.5.3.2					
IB-CentOS 6.1 Out-of-Box	InfiniBand using SLES 11 SP1	ISO Installation Only				
IB-RHEL 6.1 Out-of-Box	InfiniBand using RHEL 6.1	ISO Installation Only				
IB-SLES 11 SP1 Out-of-Box	InfiniBand using CentOS 6.1	ISO Installation Only				
W2K8-R2 WinOFED 3.0	InfiniBand using W2K8-R2 WinOFED 3.0					
IW-SL 6.1 OFED 1.5.4	iWARP using Scientific Linux 6.1 - OFED 1.5.4					

Problems Noted – December 2012

- Issues using OFED
 - Link Init Issue
 - Incorrect link speed during auto-negotiation
 - OFED Utilities
 - ibstat and ibstatus return wrong speeds
 - ibdiagnet wont read previously written topology file
- Issues using Out-of-Box Distros (CentOS, RHEL, SUSE)
 - General Setup Issues
 - ISO installations incomplete – don't install required OFA packages by default
 - Requires use of the Repositories to get all components required to support RDMA ULPs.
 - Configuration is different from OFED – for example IPoIB
 - Some ULPs not supported (NFSoRDMA, RDS, SDP)
 - IPoIB
 - Packet loss when running IPoIB in connected mode using packet sizes ≥ 8192
 - NFSoRDMA
 - RPC program not registered and cannot insert kernel modules on legacy OFED
 - Uses different port from upstream port
 - RDS
 - Did not work on any Out-of-the Box Distro
 - Big difference between upstream version and Oracle
 - SDP – Failures due to technical issues and licensing issues
 - uDAPL
 - Not working on some Distros others require installation of daplttest utilities not included on ISO

2012 Program Summary

- Automated Test Suites
 - All OFILP test scripts can now be run over VPN
- Cluster Availability
 - VPN access is granted to all OFILG Members.
 - Week long testing slots are available
- Complete heterogeneous cluster environment
 - Complete selection of IB and iWARP hardware available
 - Setup your cluster requirements in less than 30 minutes
- Distro testing as well as OFED testing available
- Support for both Windows and Linux
- Superb resources available at the UNH-IOL
 - Ethernet as well as RDMA testing available
 - Home for the OFA Training Program
 - Resident RDMA Verbs expert – Dr Robert Russell

OFA Program Details



- OFA Logo Grants: <http://iol.unh.edu/ofilglist>
- OFA Logo Program: <http://iol.unh.edu/ofilp>
- OFA Test Plan: <http://www.iol.unh.edu/services/testing/ofa/testsuites/>
- UNH-IOL: <http://www.iol.unh.edu/services/testing/ofa/>
- OFA Training Programs:
<https://www.openfabrics.org/resources/training/training-offerings.html>

Future Logo Program updates



- Distro Testing using repository updates
- RoCE Validation
 - Emulex, Mellanox, QLogic, System Fabric Works
- Interoperability testing at scale
 - 1000 nodes using Hyperion at LLNL
- OEM Interoperability validation
 - Validate new silicon and firmware releases
- IPv6 testing

OFA Logo Group

We need your support



- Planned Interop Events for 2012
 - Spring and Fall
 - Two Interop Debug events based on RC
 - Two Logo Validation events based on GA
 - Summer
 - Interop Debug Event for OFED 3.2
- The OpenFabrics Interoperability Logo Group (OFILG) needs more members
 - Just like National Public Radio
- Logo Program needs more marketing & promotion

Open Fabrics Interoperability Logo Group

