



OFED Adoption Inhibitors

within the Enterprise Data Center

Stewart Tate, tates@us.ibm.com
April 2011

Introduction



- IBM Silicon Valley Lab - Senior Technical Staff
- Almost 30 years working w/Enterprise Accounts
 - Walmart, LLBean, eBay, Borders, Macy's, Sear's
 - Citigroup, Chase, SSB
 - IRS, ATO, USAF
- Responsible for architecture/design/development
 - Hardware and Software
 - High Availability, performance, scaling

Introduction

- Currently focused on performance projects with
 - DB2 pureScale (cluster database)
 - Linux
 - AIX
 - BigData (Hadoop)
 - Linux
- IB experience for past six years
 - OFED experience past three years
 - SLES and RHEL

Introduction

- (ADD PHOTO)

Today's Enterprise Data Center



- Online Banking
 - Used daily by millions of users including cell phones
 - Citi, BofA, Chase, Wells Fargo, etc...
- Online Retail
 - Used daily by millions of users, multiple devices
 - Amazon, eBay, QVC, Apple, LLBean, etc...
- Online Entertainment
 - Rapid Growth Area requiring streaming video/audio
 - Netflix, Hulu, Fox, CNBC, ABC, NBC, CBS, etc...

Today's Enterprise Data Center



- **Continuous operation is CRITICAL to business success**
- 24 x 7 x 365 operation
- Live Upgrades / Changes to the infrastructure
- Designs allow Rapid Expansion / Contraction
- Redundant, High Availability Design

OFED addresses a number of
critical EDC requirements!

OFED Positives

- Foundational to Cloud Computing on Linux
 - HPC Clusters
 - OLTP Clusters
 - Hadoop or Data Warehousing Clusters
- Provides a High Bandwidth, Low Latency Interconnect between servers
 - Allows RDMA
 - Allows rapid data movement or large data sets

OFED Positives



- Ability to combine network fabric
 - IB Infiniband traffic
 - IP Ethernet traffic
 - SCSI Fiberchannel traffic

OFED Positives

- Numerous High Bandwidth, Low Latency Communication Options
 - IPoIB
 - uDAPL
 - SRP
 - iSER

OFED Challenges

Let's share a few examples, including a few rough edges, that could be smoother

OFED Challenges

- Example#1
 - **New SLES 10.2 System Install**
- Findings
 - Stand-alone OFED Install works GREAT!
 - Documentation Excellent!
 - Text Based Configuration Tool Perfect!

OFED Challenges

- Example#2
 - **SLES 10.2 System upgraded to SLES 10.3**
 - **OFED configuration working prior to upgrade**
- Findings
 - OFED Configuration continues to function!
 - Stand-alone OFED Install tool does not warn of distro OFED integration.
 - Stand-alone OFED does not allow additional options to be installed
 - Text Based Configuration Tool for IPoIB FAILS!
 - Documentation Invalid for distro integration!

OFED Challenges

- Example#3
 - **SLES 10.2 System upgraded to SLES 10.3**
 - **OFED not configured prior to upgrade**
- Findings
 - OFED Required Manual Configuration for IPoIB!
 - Stand-alone OFED Install tool does not warn of distro OFED integration.
 - Stand-alone OFED does not allow additional options to be installed
 - Text Based Configuration Tool for IPoIB FAILS!
 - Documentation Invalid for distro integration!

OFED Challenges

- Example#4
 - **New SLES 10.3 System Install**
- Findings
 - OFED Required Manual Configuration for IPoIB!
 - Stand-alone OFED Install tool does not warn of distro OFED integration.
 - Stand-alone OFED does not allow additional options to be installed
 - Text Based Configuration Tool for IPoIB FAILS!
 - Documentation Invalid for distro integration!

OFED Challenges

- Example#6
 - **Multi-Vendor IB switch environment issues with IPoIB**
- Findings
 - IPoIB does not function between all switches without patches
 - Various subnet managers do not support a multi-vendor IPoIB switch configuration

OFED Inhibitors

to wide adoption in the
Enterprise Data Center

Main Inhibitors

- Cost of IB .vs. Ethernet & Fiber Channel
- Fit-n-Finish within OFED
 - A few rough edges need to be polished

Inhibitor & Solution



- Inhibitor
 - Cost of IB network (card, switch, cabling) .vs.
 - Ethernet 1GB, 10GB
 - Fiber channel 2GB, 4GB, 8GB
- Solution
 - Promote Cost Benefits of Combined Fabrics
 - QDR IB + Faster Ethernet + Faster Fiber channel SAN over IB
 - Fewer switches
 - Less Cabling
 - Lower Labor Costs
 - Promote Low Latency and Higher Bandwidth
 - IB, uDAPL, IBverbs
 - Faster networks (IB+IP+FC) for less than separate networks

Inhibitor & Solution

- Inhibitor
 - Fit-n-Finish of OFED .vs.
 - Ethernet 1GB, 10GB
 - Fiber channel 2GB, 4GB, 8GB
- Solution
 - Polish the rough edges
 - Distro Integration as good or better then stand alone package
 - Install / Config GUI for install, upgrade, options, validation
 - Integrated, easy to locate, documentation
 - Improved vendor hardware / software integration
 - Strict Quality Control Standards

Thank you!