



COMPONENTS

- Driver & Access Layers
- OSU MVAPICH and Open MPI*
- IP-over-IB
- SDP
- SRP Initiator
- iSER Initiator*
- uDAPL*
- RDS
- Subnet Manager (OpenSM)
- Installation, Administration and Diagnostics Tools
- Performance test suites
- * Currently not supported by Mellanox

DEVICES SUPPORT

- InfiniHost HCA Silicon
- InfiniHost III Ex HCA Silicon
- InfiniHost III Lx HCA Silicon
- ConnectX HCA Silicon
- Memory & Memory-free HCA Cards
- InfiniScale Switch Silicon
- InfiniScale III Switch Silicon

BENEFITS

- Single software stack that operates across all available InfiniBand devices and configurations such as mem-free, DDR/SDR, PCI-X, and PCI Express modes
- Support for HPC applications for scientific research, oil and gas exploration, car crash tests, bench marking etc. E.g., Fluent, LS-Dyna
- Support for EDC applications such as Oracle 10g RAC, IBM DB2, SAP, and financial applications
- Support for traditional IP and Sockets based applications leveraging the benefits of RDMA
- Support for high performance block storage applications utilizing RDMA benefits
- Support for cluster file systems for high performance and virtualized file system access
- Support for cluster management tools

OFED (OpenFabrics Enterprise Distribution)

High performance server and storage connectivity software for field-proven RDMA and Transport Offload hardware solutions

Linux Edition

Overview

Use of clustered commodity servers, in lieu of traditional supercomputers and mainframes, offers tremendous price/performance benefits and unparalleled flexibility in deployment and long-term maintenance. To enable distributed computing transparently, high performance computing (HPC) applications require the highest bandwidth and lowest possible latency. In enterprise data center (EDC) applications, these requirements are compounded with the need to support a large interoperable ecosystem of networking, virtualization, storage, and other applications and interfaces.

The OFED from OpenFabrics alliance (www.openfabrics.org) has been hardened through collaborative development and testing by all major InfiniBand vendors. OFED is supported by Mellanox and major InfiniBand vendors to enable OEMs and System Integrators to meet the needs of HPC and EDC applications.

Attain Higher Bandwidth and Lower Latency

For HPC applications, OFED offers Message Passing Interface (MPI) implementation from Ohio State University (OSU) and Open MPI (www.open-mpi.org). Bandwidth results in excess of 1400MB/s and application latencies lower than 3 microseconds have been achieved. They include features to enable scaling to large clusters while improving on memory usage and latency related efficiencies. They are optimized and tested for use with Fortran, C and C++ compilers from popular vendors.

For traditional TCP/IP and sockets-based applications, OFED includes the field proven implementation of IP-over-IB (as defined by the IETF, www.IETF.org), enabling IP-based applications to work seamlessly over InfiniBand and perform at levels more than three times higher than Ethernet. It also includes the IBTA (www.IBTA.org) defined Sockets Direct Protocol (SDP) enabling traditional TCP/IP sockets-based applications to capitalize

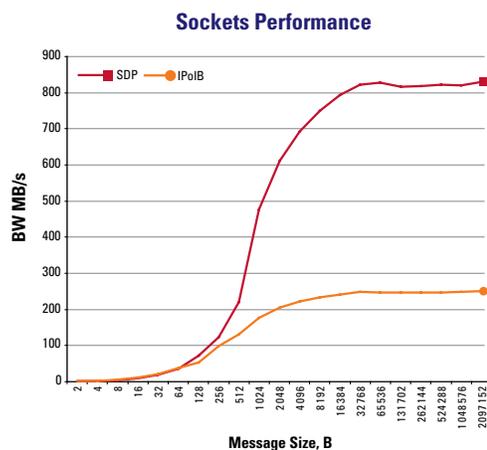


Figure 1: SDP enables 3x performance improvement over Ethernet & TCP/IP

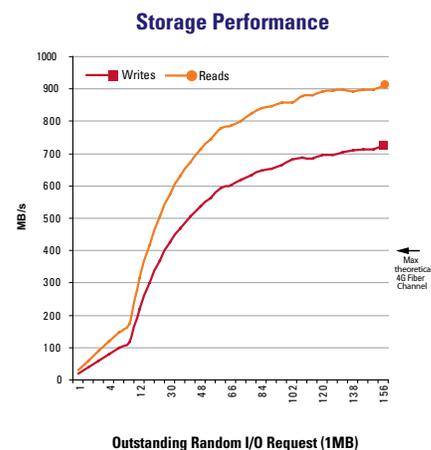


Figure 2: SRP performance for random Read/Write has 2x performance improvement over 4G Fiber Channel

OFED (OpenFabrics Enterprise Distribution)

BENEFITS

Sockets Layer

- SDP and IP-over-IB component enable TCP/IP and sockets-based applications to interface seamlessly to and benefit from InfiniBand transport

Access Layer

- Supports the OpenFabrics defined Verbs API at the user and kernel levels. User level verbs allow protocols MPI and other applications to interface to Mellanox InfiniBand hardware. Kernel level verbs allow protocols like SDP, SRP and IP-over-IB to interface to Mellanox InfiniBand hardware

SCSI Mid Layer

- The SCSI Mid Layer interface enables SCSI based block storage and management applications to interface with the SRP Initiator component and the Mellanox InfiniBand hardware

on the RDMA and transport offload capabilities of InfiniBand.

To enable traditional SCSI and iSCSI-based storage applications to enjoy similar RDMA performance benefits, OFED includes the SCSI over RDMA Protocol (SRP) initiator and the iSCSI RDMA Protocol (iSER) that interoperate with various target components available in the industry. The SRP initiator, for example, interoperates with the SRP target reference implementation from Mellanox, which is open source and available freely from www.openfabrics.org. SRP solutions over InfiniBand have been proven to deliver impressive 910MB/s (random read) and 725MB/s (random write) I/O storage performance with 1MB blocks.

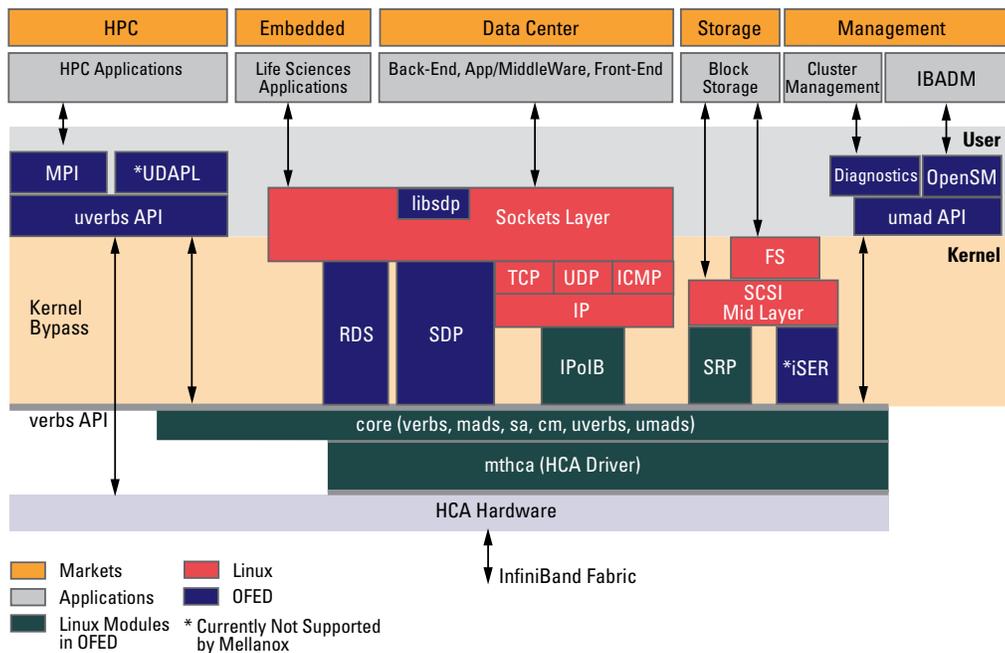
OFED is interoperable across multi-vendor vendor solutions

As a founding member of OpenFabrics, previously known as OpenIB, Mellanox is driving interoperability of the OpenFabrics software across different vendor solutions. OFED, running on servers, interoperates with InfiniBand switches, gateways and storage targets and value added software provided by vendors in such systems. OFED is also planned to be distributed as part of major Linux OS distributions. OFED supports all popular CPU platforms.

erability of the OpenFabrics software across different vendor solutions. OFED, running on servers, interoperates with InfiniBand switches, gateways and storage targets and value added software provided by vendors in such systems. OFED is also planned to be distributed as part of major Linux OS distributions. OFED supports all popular CPU platforms.

Enabling a Large Software Ecosystem

Through its collaboration with the OpenFabrics community, support for multiple industry-standard interfaces that enable off-the-shelf applications to work with InfiniBand, and support for popular OS and CPU platforms, Mellanox has created a large, interoperable software ecosystem that benefits the HPC and EDC markets.



Platform Supplier Products and Support	
RedHat	AS 4.0 – Update 2 and Update 3, Fedora Core 4
Novell	SLES 9 – SP3, SUSE 10 Pro, SLES 10
Intel	x86, IA64, EM64T, EM64T x86_64, PCI-X and PCIe platforms
AMD	AMD64 and AMD64-Ex Opteron, PCI-X and PCIe platforms



2900 Stender Way, Santa Clara, CA 95054
 Tel: 408-970-3400 • Fax: 408-970-3403
www.mellanox.com