



InfiniBand Gold (IBGold) Collection

Product Brief

- Linux Edition -



HIGH PERFORMANCE INTERCONNECT SOFTWARE FOR HPC AND ENTERPRISE DATA CENTERS

IBGold Components

- Driver & Access Layers
- MPI
- IP-over-IB
- SDP
- SRP Initiator
- SRP Target (separate package)
- uDAPL
- kDAPL
- Subnet Manager (OpenSM)
- Installation and Management Tools

InfiniBand Devices Support

- InfiniHost HCA Silicon
- InfiniHost III Ex HCA Silicon
- InfiniHost III Lx 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.
- Support for EDC applications such as Oracle RAC, IBM DB2, SAP, and financial applications
- Support for traditional IP and Sockets based applications, both over TCP/IP, as well utilizing RDMA and kernel bypass modes for higher performance, lower CPU utilization and lower latency
- 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

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, 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 IBGold Collection from Mellanox Technologies is designed, packaged and supported to enable OEMs to meet the needs of HPC and EDC applications.

Attain Higher Bandwidth & Lower Latency

For HPC applications, IBGold offers the popular and well-deployed Message Passing Interface (MPI) implementation from Ohio State University (OSU). Bandwidth results in excess of 1400MB/s and application latencies lower than 3 microseconds have been achieved. Mellanox has added significant enhancements to its MPI implementation to enable scaling to large clusters while improving on memory usage and latency related efficiencies. The IBGold MPI implementation is optimized and tested for use with popular Fortran, C and C++ compilers from Intel, PGI and others.

For traditional TCP/IP and sockets-based applications, IBGold offers a robust and field proven implementation of IP-over-IB to enable IP-based applications to work seamlessly over InfiniBand and perform at levels higher than Ethernet. IBGold also includes an InfiniBand optimized implementation of the IBTA defined Sockets Direct Protocol (SDP) for enabling traditional TCP/IP sockets-based applications to capitalize on the RDMA (zero-copy) and transport offload capabilities of InfiniBand.

To enable traditional SCSI-based storage applications to enjoy similar RDMA performance benefits, IBGold provides SCSI over RDMA Protocol (SRP) initiator and target components. Using Oracle's Orion benchmark tests, the IBGold SRP solution has delivered impressive 910MB/s (random read) and 725MB/s (random write) I/O storage performance with 1MB blocks. IBGold extends RDMA capabilities to Network File System (NFS) applications as well using NFS-RDMA solutions that will be available in upcoming releases.

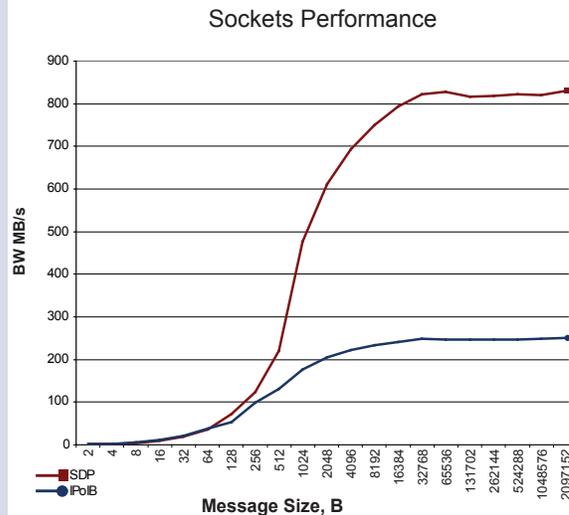


Figure 1: SDP enables 3x performance improvement over Ethernet & TCP/IP

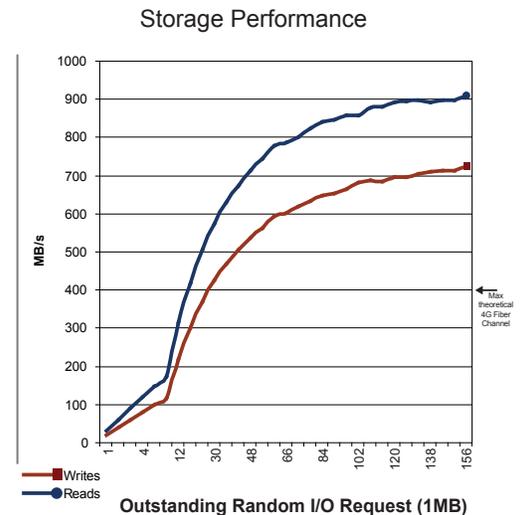


Figure 2: SRP performance for random Read/Write has 2x performance improvement over 4G Fiber Channel

Enabling Easy OEM Adoption & Deployment

- Simple open source licensing terms facilitate easy adoption and integration of IBGold into the OEM's value added software (excludes SRP Target component)
- Comprehensive base layer support for HPC and EDC applications
- Installation, debug and management tools reduces the OEM and its end users' total cost of ownership
- ISV partners offer pre-integrated IBGold vertical solutions to further accelerate time to market

Installation and Management Tools

- Scripts for quick and easy installation of IBGold components
- Binary RPMs and source tarballs for advanced installations and broader platform support
- Pre-boot Execution Environment (PXE) over InfiniBand to facilitate network boot and remote installations
- Command line based debug and cluster verification tools designed for ensuring that cluster designs are done right the first time, and to accelerate cluster bring up and debug time
- Firmware generation and burning tools

Vertical Solutions Supported

- ISV applications such as those for cluster file system, storage or scientific research, are pre-integrated and tested with IBGold
- IBM DB2, Oracle
- TerraScale TerraGrid scalable storage and database solutions
- Cluster File System Inc. Lustre
- NPACI Rocks Cluster Package
- Fluent 6.2
- Virtual Iron Virtualization Software

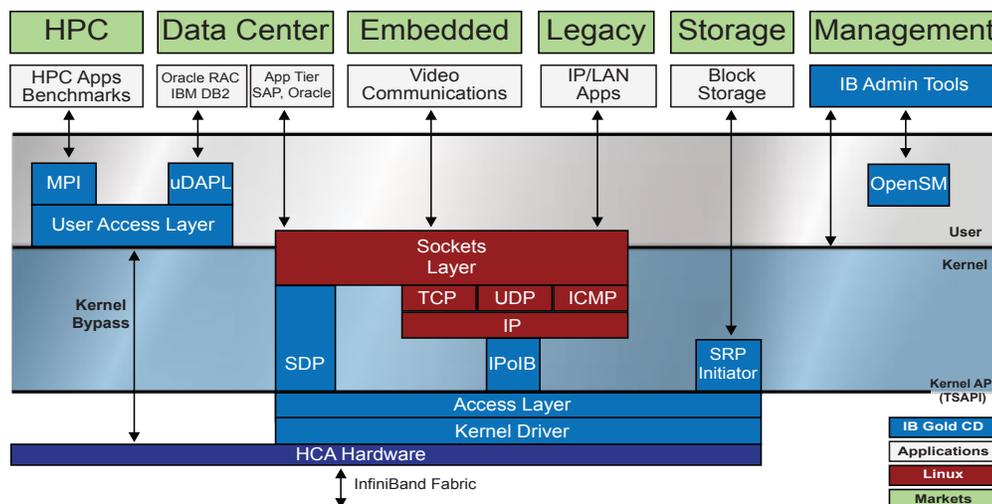
Support for a Large and Interoperable Software Ecosystem

Through its collaboration with OpenIB, 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.

As a founding member of OpenIB—the organization chartered with creating interoperable Linux kernel-level stacks for InfiniBand—Mellanox actively contributes software components to this open-source community. In addition, Mellanox HCA solutions are used as the gold standard for hardware in OpenIB, and future IBGold versions will include qualified and released versions of the OpenIB Gen2 InfiniBand stack components.

While HPC applications mostly use the Message Passing Interface (MPI), storage, embedded and IP applications prevalent in the EDC require multiple interfaces to be supported, each with its unique performance optimization and application programming interface capabilities to provide the highest benefits to the serviced user level applications.

Interface Supported	Benefits
Sockets Layer	SDP and IP-over-IB components enable TCP/IP and sockets-based applications to interface seamlessly to and benefit from InfiniBand transport
Access Layer	Supports implementations of IBTA defined Verbs API or (VAPI) at the user and kernel levels to allow MPI and other applications to interface to Mellanox InfiniBand hardware (Next generation of IBGold supports the OpenIB defined driver access layer for maximum upper layer protocol and application level interoperability)
uDAPL	Supports user level and kernel level DAPL interfaces for enabling file system, database and block storage applications to interface to Mellanox InfiniBand hardware



Support for Popular and High Performance OS and CPU Platforms

Platform Supplier	Products and Support
RedHat	Latest 32-bit and 64-bit Linux OS distributions - Advanced Server and Fedora Core 3
Novell	Latest SUSE Linux Enterprise Edition or SLES and SUSE Linux Professional
Intel	32-bit and 64-bit CPUs including the Intel IA32, IA64, EM64T based PCI-X and PCIe platforms
AMD	32-bit Athlon, and 64-bit Opteron PCI-X and PCIe platforms



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