Performance Enhanced OpenMPI/OFED

Mercury's Open MPI on Serial RapidIO®

- Enables higher performance communication using enhanced MPI allocation
- Provides MPI portability and code compatibility across high-speed fabrics and processor architectures
- Supports zero-copy RDMA message-passing with application buffers
- Higher performance shared memory communication with optimized memory copy
- High throughput and low latency on RapidIO with minimal overhead

Performance Enhanced OpenMPI/OFED from Mercury Computer Systems optimizes Open MPI (the open-source high-performance computing library) with Mercury's OFED provider software module for serial RapidIO. Mercury has enhanced the OFED provider software for better throughput and lower latency, based on years of experience in delivering high-performance middleware and serial RapidIO solutions.

OpenMPI/OFED implements the open-standard MPI allocation functions using DMA-optimized buffer allocation methods, resulting in significantly improved application performance. These enhancements combine the productivity and ease-of-use of the MPI API, designed for data-intensive computing, with the speed of an optimized low-latency, high-bandwidth switch fabric, designed for embedded computing.

Mercury-Optimized Middleware with Open-Source Portability and Compatibility

OpenMPI/OFED provides a release of the OpenFabrics Enterprise Distribution (OFED) to serial RapidIO, enabling, for example, high-performance MPI applications to run on Mercury multicomputers. OpenMPI/OFED runs directly on the high-speed fabric, offering an optimized open-source alternative for data-plane communications. By providing an optimized OFED library for serial RapidIO, Mercury has made Open MPI suitable for dense image and signal processing solutions utilizing PPC or Intel POET™-based processing modules.

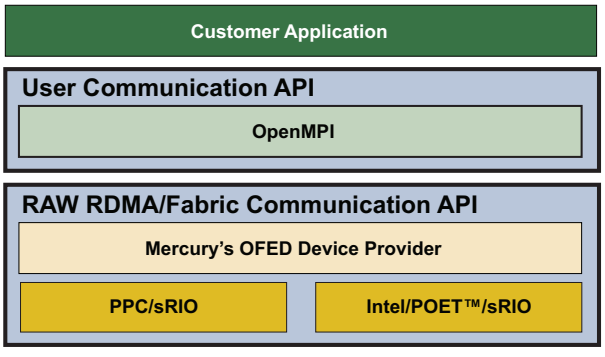

***Figure 1**. OpenMPI/OFED Architecture*

MPI applications currently running on InfiniBand® fabrics can be re-compiled without modification to run in a serial RapidIO embedded environment. Compatibility is realized across all fabrics with the high-performance middleware portability of OFED. Overall software savings are realized through application "re-use" with open-standard software, software-compatible ports, and taking advantage of the high performance of embedded computing.

OpenMPI/OFED incorporates Mercury's more than 15 years of experience in optimized middleware with a commitment to bringing this experience to open solutions. Mercury's Interprocessor-Communication System (ICS) was the industry's first multicomputer operating environment, providing both high-speed data movement and synchronization. The Parallel Acceleration System (PAS™) was the industry's first middleware that enabled high-speed, low-latency applications to use complex data reorganization patterns.

Features and Benefits

- Application portability among multiple systems using serial RapidIO, Ethernet, or other high-speed fabrics
- Integration with ICS for ultimate inner loop performance
- Alternative MPI implementations or OFED-dependent software supported as an optional service
- Migration between PPC-based and x86-based processing modules
- Low latency
- High throughput

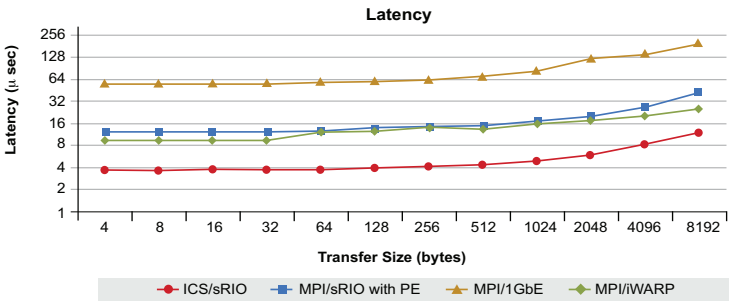

***Figure 2**. Minimal transfer startup overhead using high-level open-source middleware*

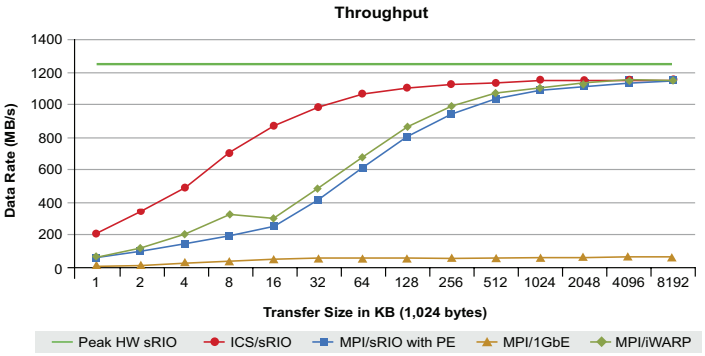

***Figure 3**. Full sRIO throughput using high-level open-source middleware*

OFED

The OpenFabrics Enterprise Distribution (OFED) implements channel I/O and enables RDMA on fabrics and networks in Linux® and Windows® environments. The software stack includes middleware and drivers for 10 Gigabit Ethernet and 10/20/40 Gigabit InfiniBand® interconnects. The RDMA services include MPI in clusters and systems, legacy IP networking, uDAPL, NFS, RDS, iSCSI (iSER), and SCSI Remote (SRP) protocols, as well as many parallel and database file systems, such as Lustre® and Oracle®. OFED delivers the highest achievable data-transfer rates and the lowest achievable latencies for the full range of virtualization and critical solution applications, cloud computing, and simulation/modeling in high-performance computing (HPC).

OpenFabrics Alliance

The OpenFabrics Alliance (OFA) is an industry standard consortium that develops, tests, licenses, and distributes cross-platform, open-source middleware for high-performance applications, low-latency Ethernet and InfiniBand networks, wire-speed networking, and microsecond latencies. OpenFabrics software is used in critical applications that require highly efficient networks, storage connectivity, and parallel multiprocessor computing. The software provides HPC sites and enterprise solution centers with flexibility and investment protection, as computing evolves toward applications that require extreme speed, massive scalability, and utility-class reliability.

More than 40 percent of the 100 top-performing HPC systems and as many as 60 percent of all new HPC installations worldwide use OpenFabrics software as an Enterprise Distribution (OFED) for parallel multicomputing, low-latency interconnects, and/or file-system operations. Mercury is a member of the OpenFabrics Alliance (OFA).

Learn More

Mercury Middleware
www.mc.com/products/software/middleware.aspx

OpenFabrics Alliance (OFA)
www.openfabrics.org

Message Passing Interface (MPI)
www.mpi-forum.org

Open MPI
www.open-mpi.org

NetPIPE benchmark
www.scl.ameslab.gov/netpipe

RDMA over TCP/IP – iWARP
www.rdmaconsortium.org/

System Requirements

Ensemble™ Series platforms with MultiCore Plus software on

Linux OS

864x PowePC®	Serial RapidIO
x86 Intel	Serial RapidIO

Ordering Information

Contact your Mercury representative for a free product demonstration and ordering information.

 2507.01E-0411-DS-openmpi

MERCURY COMPUTER SYSTEMS™

Corporate Headquarters
201 Riverneck Road
Chelmsford, MA 01824-2820 USA
+1 (978) 967-1401 • +1 (866) 627-6951
Fax +1 (978) 256-3599
www.mc.com

Europe
Mercury Computer Systems, Ltd.
Unit 1 - Easter Park, Benyon Road • Silchester, Reading • RG7 2PQ UNITED KINGDOM
+ 44 0 1189 702050 • Fax + 44 0 1189 702321

Asia
Nihon Mercury Computer Systems K.K.
No. 2 Gotanda Fujikoshi Bldg. 4F • 5-23-1 Higashi Gotanda • Shinagawa-ku, Tokyo 141-0022 JAPAN
+81 3 3473 0140 • Fax +81 3 3473 0141

Challenges Drive Innovation®