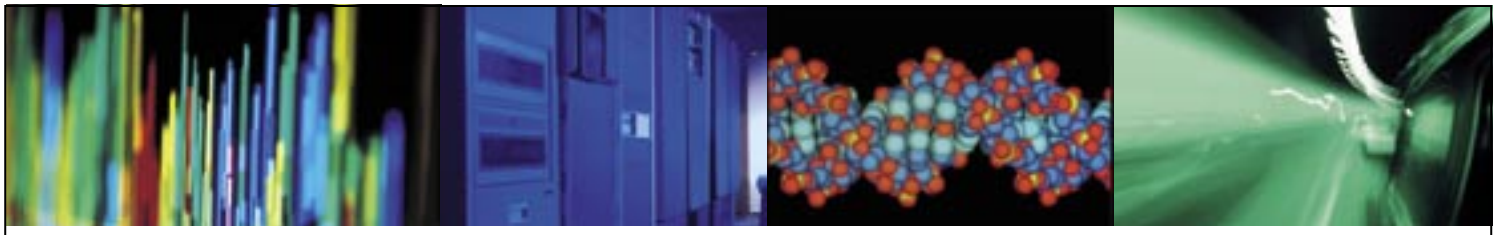

PathScale™

PATHSCALE INFINIPATH™ HTX™ ADAPTER
LOW-LATENCY CLUSTER INTERCONNECT FOR INFINIBAND™

The PathScale InfiniPath HTX Adapter provides the industry's lowest latency cluster interconnect for MPI applications. The adapter directly attaches to AMD Opteron™ processor-based servers through a standard HyperTransport™ HTX slot on an Opteron motherboard. To connect to other nodes in the cluster, the adapter utilizes standard InfiniBand switches and cabling. These standards enable one to build Linux clusters with the best combination of low-latency and high bandwidth to achieve massive scalability, while reducing costs and complexity. The result is an unparalleled 1.5 µs MPI latency and 1.8 GB/sec of bi-directional bandwidth.

INDUSTRY STANDARDS REDUCE COST AND COMPLEXITY

The PathScale InfiniPath HTX Adapter supports a rich combination of open standards to achieve leading performance, while taking advantage of economies of scale. The InfiniPath HTX Adapter connects directly to the Opteron CPU via an open standard HyperTransport HTX slot and is compliant with the InfiniBand link and physical layers, providing end-to-end compatibility. Any InfiniBand-compliant switch can be used as the external switching fabric for the InfiniPath interconnect. The PathScale InfiniPath Software Stack includes an optimized MPICH compatible Message Passing Interface (MPI) library and supports the most recent standard Linux distributions from SuSE™ and Red Hat™.

LOWEST LATENCY

The PathScale InfiniPath HTX Adapter creates a direct connection between an AMD Opteron-processor based server and the InfiniBand switch with 1.5 µs of MPI latency. By eliminating the excess latency typically found in current InfiniBand adapters, communications wait time is reduced and processors spend more time computing.

HIGH BANDWIDTH

The InfiniPath HTX Adapter uses the HyperTransport interface at a peak rate of 6.4GB/s and InfiniBand at a peak rate of 2GB/s of bi-directional network bandwidth. The MPI software stack enables applications to deliver 1.8GB/s of bi-directional bandwidth. There are no bandwidth compromises.

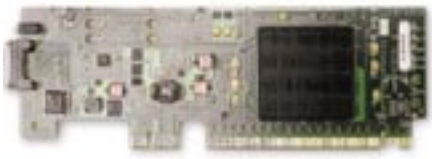

UNMATCHED SCALABILITY

While 0-byte latency and peak bandwidth for long messages are key metrics often reported for an interconnect, the scalabilty of most real applications is determined by the performance of the interconnect on small to medium-sized packets. The PathScale InfiniPath interconnect excels in this regime, achieving the crucial 1/2 peak streaming bandwidth point at packet sizes significantly below both traditional InfiniBand and proprietary cluster interconnects.

Communications overhead is reduced for real applications, with the result that superior scalability is achieved. With InfiniPath, applications can now efficiently scale to thousands of nodes.

INFINIPATH SOFTWARE STACK

Developed specifically for HPC Clusters, the PathScale InfiniPath software drivers and libraries are optimized for message passing, with streamlined protocols (MPI and IP) encapsulated on top of standard InfiniBand wire protocols and fabric. Most of the protocol processing is handled directly by the Opteron processor, not a slower embedded processor on the adapter. This design allows for easy software upgrades, and a simpler, more reliable design.

PATHSCALE INFINIPATH HTX ADAPTER PRODUCT BENEFITS

LOWEST MPI LATENCY

- 1.5 µs latency for 0-byte messages through an InfiniBand switch
- Performance improves with increased CPU and memory clock speed

IMPLEMENTED ON INDUSTRY STANDARDS

- HyperTransport 1.03
- InfiniBand 1.1 Switches & Cabling
- MPI 1.0 with MPICH 1.2.6; 2.0 future

HIGHEST SUSTAINABLE BANDWIDTH

- 1.8 GB/s bi-directional bandwidth
- N½ streaming message size of 595 bytes

EASY TO MANAGE AND SUPPORT

- Standard InfiniBand fabric management
- No costly hardware, software, or operating system upgrades required
- Supported with SuSE, Red Hat and Fedora Core Linux

PATHSCALE: THE FOUNDATION FOR PERFORMANCE AND RELIABILITY

PathScale is dedicated to accelerating application performance on Linux clusters. The combination of the PathScale InfiniPath Interconnect, the industry leading PathScale EKOPath™ Compiler Suite and the PathScale OptiPath™ MPI Acceleration Tools support the company mission to accelerate the development and deployment of 64-bit applications on Linux clusters.

AMD 64 Opteron

PATHSCALE INFINIPATH™ HTX™ ADAPTER

HYPERTRANSPORT INTERFACE
- HT v1.0.3 compliant
- HTX slot compliant
- Low profile design, single HyperTransport port
- 6.4 GB/s of processor bandwidth

CONNECTIVITY
- Single 4X (10 Gb/s) port – Copper
- Compatible with InfiniBand switches from Infinicon™, Mellanox™, TopSpin™ and Voltaire™

PATHSCALE HOST DRIVER/UPPER LEVEL PROTOCOL (ULP) SUPPORT
- MPICH version 1.2.6 with MPI 2.0 ROMIO I/O
- TCP, NFS, UDP, SOCKETS through Ethernet driver emulation
- Optimized MPI protocol stack supplied
- 32 and 64-bit application ready

INFINIBAND INTERFACES AND SPECIFICATIONS
- InfiniBand link layer compliant
- Uses standard IBTA 1.1 compliant fabric and cables, Link layer compatible
- Configurable MTU size (4096 maximum)

MANAGEMENT SUPPORT
- Includes an SMA 1.1 compliant subnet management agent
- Interoperable with management solutions from Infinicon, Topspin and Voltaire

OPERATING ENVIRONMENTS
- Supported Linux Versions
  - Fedora Core 2 with 2.6 kernel
  - SuSE Linux 9.1 & SLES 9
  - Red Hat Enterprise Linux 3.x

PATHSCALE INFINIPATH ADAPTER SPECIFICATIONS
- Maximum Power Consumption: 9 Watts
- Typical Power Consumption: 7 Watts
- Operating Temperature: 10 to 45C at 0-3km
  –30 to 60C (Non-operating)
- Humidity:
  20% to 80% (Non-condensing, Operating)
  5% to 90% (Non-operating)

PathScale™

Accelerating Cluster Performance

PATHSCALE, INC.
2071 STIERLIN COURT, SUITE 200
MOUNTAIN VIEW, CA 94043 USA

TEL 650.934.8100
FAX 650.428.1969
PATHSCALE.COM
SALES@PATHSCALE.COM

AMD 64 Opteron