


# NVIDIA GPUDirect™ Technology – Accelerating GPU-based Systems

The rapid increase in the performance of graphics hardware, coupled with recent improvements in its programmability, has made graphic accelerators a compelling platform for computationally-demanding tasks in a wide variety of application domains. Due to the great computational power of the GPU, the GPU-to-GPU method has proven valuable in various areas of science and technology.

GPU-based clusters are being used to perform compute-intensive tasks, like finite element computations, computational fluids dynamics, Monte-Carlo simulations, etc. Several of the world-leading supercomputers are using GPUs in order to achieve the desired performance. Since GPUs provide high core count and floating-point operations capabilities, high-speed InfiniBand networking is required to connect between the platforms in order to provide high throughput and the lowest latency for GPU-to-GPU communications.


While GPUs have been shown to provide worthwhile performance acceleration yielding both price/performance and power/performance benefits, several areas of GPU-based clusters could be improved in order to provide higher performance and efficiency. The main performance issue with deploying clusters consisting of multi-GPU nodes involves the interaction between the GPUs, or the GPU-to-GPU communication model. Prior to the GPU-Direct technology, any communication between GPUs had to involve the host CPU and required buffer copies. The GPU communication model required the CPU to initiate and manage memory transfers between the GPUs and the InfiniBand network. Each GPU-to-GPU communication had to follow the following steps:

1. The GPU writes data to a host memory dedicated to the GPU
2. The host CPU copies the data from the GPU dedicated host memory to host memory available for the InfiniBand devices for RDMA communications
3. The InfiniBand device reads the data and sends it to the remote node




CPU involvement in the GPU communications and the need for a buffer copy has created bottlenecks in the system, slowing the data delivery between the GPUs.

The new GPUDirect technology from NVIDIA and Mellanox enables NVIDIA Tesla and Fermi GPUs to communicate faster by eliminating the need for CPU involvement in the communication loop and the need for the buffer copy. The result is increased overall system performance and efficiency by reducing the GPU-to-GPU communication time by 30%. NVIDIA GPUDirect provides a new interface between the GPU and the Mellanox InfiniBand adapters and enables both devices to share the same system memory.



The performance gain for high-performance applications depends on the amount of GPU communication being used. Applications that utilize parallel execution can see performance gain or productivity increase of up to 42%. All applications will show performance and efficiency improvements with Mellanox InfiniBand adapters and NVIDIA GPUDirect technology. Mellanox InfiniBand adapters with NVIDIA GPUDirect is an essential technology for GPU-based systems. The combined solution delivers the capability to maximize the performance capability of the GPUs and the overall system productivity, delivering the highest return-on-investment.

## ConnectX<sup>®</sup> 2



350 Oakmead Parkway  
Sunnyvale, CA 94085

Tel: 408-970-3400 • Fax: 408-970-3403

[www.mellanox.com](http://www.mellanox.com)

© Copyright 2010, Mellanox Technologies. All rights reserved. Preliminary information. Subject to change without notice. Mellanox, BridgeX, ConnectX, InfiniBlast, InfiniBridge, InfiniHost, InfiniRISC, InfiniScale, InfiniPCI, PhyX and Virtual Protocol Interconnect are registered trademarks of Mellanox Technologies, Ltd. CORE-Direct and FabricIT are trademarks of Mellanox Technologies, Ltd. All other trademarks are property of their respective owners.