



Electronic Design Automation (EDA)

Design automation systems reduce risks and speed time to market by enabling engineers to collect vital information about how products will perform. EDA simulations often involve 3D modeling, fluid dynamics, and other compute-intensive processes that require high-performance computing (HPC) data center solutions.

HPC data centers now include clusters of multi-core, multi-processor servers and high-speed, high-capacity storage systems. In these environments, moving data rapidly among servers and from servers to storage and back is crucial for delivering the highest performance and productivity. Mellanox InfiniBand adapters, switches, cables and software products maximize data center performance and productivity by delivering industry-leading bandwidth and the lowest latency.

THE MELLANOX SOLUTION

Mellanox I/O adapters, switches, and gateways enable data center managers to gain maximum leverage from server and storage investments by delivering up to 200Gb/s per port of reliable and lossless connectivity with sub-600 ns latency. With this level of performance, processing models and simulations at top speed ensures minimal wait times for design engineers.

For IT managers, Mellanox solutions reduce infrastructure costs and complexity by slashing the number of required I/O adapters, reducing management overhead and cutting power consumption. They also enable storage and server traffic to consolidate on a single unified fabric.

Mellanox solutions are built on open standards delivering exceptional performance and business value. Only Mellanox offers backwards and forwards compatibility, which future proofs your networking investment. What's more, Mellanox is the only interconnect that is built on offload architecture, freeing more CPU cycles for application processing with no CPU frequency sensitivity. MPI collectives are managed and executed from the network fabric and is the only interconnect that offers in-

KEY BENEFITS

- The world's fastest interconnect supporting up to 200Gb/s per port
- In-network co-processing capabilities and advanced offload engines
- Sub-600 ns latency
- Lossless packet transmission
- Converged data center fabric reduces costs and power requirements
- Support for GPUDirect RDMA and GPUDirect ASYNC