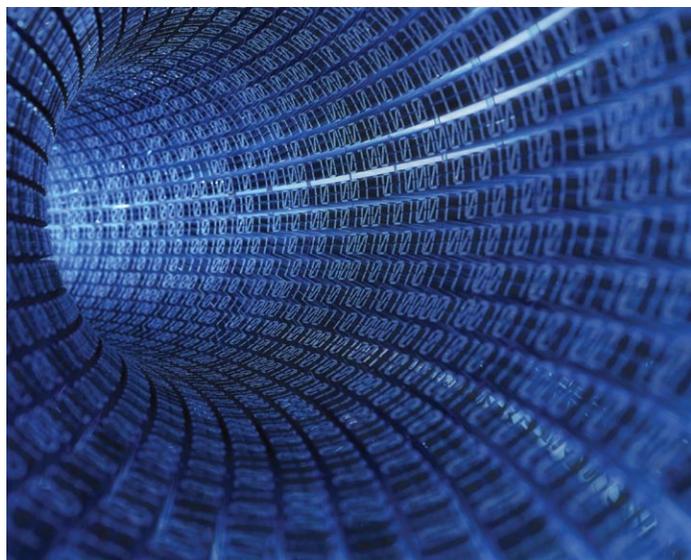


IBM HPC Insights

A regular communiqué from the IBM Deep Computing team

Achieving scalable HPC performance with IBM System Cluster 1350

Mellanox ConnectX InfiniBand adapters increase IBM System Cluster 1350 system efficiency and enhance HPC scalability, delivering outstanding performance and a strong return on hardware investments



The addition of more processing cores per node and more nodes per cluster offers the potential of using high-performance computing (HPC) for solving even larger and more complex problems. Yet the full value of many-core systems and large-scale clusters can be realized only by reducing bottlenecks within and between nodes. Using Mellanox® ConnectX® adapters with IBM® System Cluster 1350 solutions can help minimize bottlenecks. ConnectX adapters capitalize on the throughput and low latency of InfiniBand while offloading communications from the processor to the network, enabling processors to focus on application work.

OFFLOADING PROCESSOR TASKS BOOSTS SYSTEM EFFICIENCY

Multi-socket, multi-core systems used for large-scale HPC clusters might have high theoretical floating-point operations per second (FLOPS) values. But many systems—including those used in clusters on the TOP500 super-computer list—fall far short of that value in real-world use. Some clusters that use Ethernet connections between nodes achieve only 50 percent system efficiency as processors sit idle, waiting for data, or are forced to devote cycles to networking communications rather than focusing on application work.

InfiniBand adapters can help increase system efficiency. With higher bandwidth and lower latency than Ethernet, InfiniBand helps move data to and from nodes rapidly, reducing the idle time of processors. Mellanox ConnectX adapters are among the top-performing InfiniBand solutions: ConnectX adapters currently are the only native Remote Direct Memory Access (RDMA) InfiniBand solutions that move data between nodes at 40 Gb/s with latency as low as one μ sec—about the same latency as data movement between a processor and its memory.

ConnectX adapters also free up processing cycles for application work by offloading networking communications from the processor to the adapter. While Ethernet connections require the processor to devote cycles to TCP processing, and other InfiniBand solutions require processors to use cycles for RDMA and other tasks, ConnectX adapters handle all these communications within the adapter hardware. As a result, the processor can gain back additional cycles for application processing.

The high throughput, low latency and offloading capabilities offered by Mellanox adapters can help to boost system efficiency to up to 96 percent, as demonstrated by clusters on the TOP500 list. With greater system efficiency, organizations can maximize the overall performance of their cluster.

CONSOLIDATING COMPONENTS REDUCES COST AND COMPLEXITY

Using an IBM System Cluster 1350 solution with Mellanox ConnectX adapters also can help consolidate components, reducing the power, cooling and real estate required for HPC. Cluster 1350 solutions that use BladeCenter® or iDataPlex™ servers increase computing density by packing multi-socket, multi-core processing architectures into small enclosures. ConnectX adapters consolidate I/O by facilitating convergence of networking, storage, compute and management connectivity in a single wire. Using fewer cables and switches can help cut costs, save space, decrease IT complexity and reduce potential points of failure.

REACHING NEXT-LEVEL SYSTEM EFFICIENCY ACCELERATES RUN TIMES

The recent introduction of CORE-Direct™ technology by Mellanox can help ensure high system efficiency even as more cores and nodes are added to HPC clusters. Developed through a [joint project with the U.S. Department](#)



IBM HPC Insights

A regular communiqué from the IBM Deep Computing team

of Energy and the Oak Ridge National Laboratory, *CORE-Direct* focuses on the communications between nodes for data broadcast, data collection and global application synchronization—communications that can be key limiting factors in achieving parallel application efficiency and scalability for scientific simulation.

By offloading those communications from the processor to the adapter, *CORE-Direct* technology enables overlapping communications and computation so processors can perform application processing without interruptions. It also reduces the noise within systems that keeps the processor from focusing on application processing and slows overall cluster performance. Organizations using Mellanox ConnectX-2 adapters with *CORE-Direct* technology could see run times reduced by half.

INCREASING EFFICIENCY AND CONSOLIDATING COMPONENTS MAXIMIZES ROI

Increasing system efficiency and consolidating I/O are essential for making the most of HPC expenditures. When each processor devotes more cycles to application processing, organizations can solve large, complex problems while controlling the number of nodes they need to purchase, power, cool and house as part of their HPC solution. At the same time, consolidating multiple functions in a single adapter and cable can help organizations reduce I/O expenditures and further decrease ongoing power consumption. By integrating the Mellanox ConnectX family of adapters into IBM System Cluster 1350 solutions, organizations can maximize the return on their HPC investments. 

For more information about IBM System Cluster 1350 solutions, visit: ibm.com/systems/clusters/hardware/1350/index.html

To learn more about Mellanox ConnectX InfiniBand products, visit: mellanox.com

The HPC Advisory Council has published multiple application case studies; for more information, visit: www.hpcadvisorycouncil.com