



NVIDIA QUANTUM INFINIBAND ACCELERATES BIOSCIENCE

Accelerated Performance Across the Data Center

Bioscience is a field of study that seeks to develop solutions that improve our quality of life—from precise medical instruments and pharmaceuticals to digital health applications and life-saving therapies. The benefits can be realized at a global scale.

In this diverse field, researchers fuse traditional simulations with deep learning, big data analytics, and medical instruments that stream data to a central supercomputer, often composed of millions of processing cores. Today's generation of bioscience research demands high performance networks that enable ultra-fast processing of high-resolution simulations, extreme-size datasets, and highly parallelized algorithms that can exchange information in real time.

Advanced Research Needs Advanced Networking

Complex workloads demand ultra-fast processing of high-resolution simulations, extreme-size datasets, and highly parallelized algorithms. As these computing requirements continue to grow, NVIDIA Quantum InfiniBand provides the dramatic leap in performance needed to achieve unmatched performance in HPC, AI, and hyperscale cloud infrastructures with less cost and complexity.

NVIDIA Quantum InfiniBand delivers high-speed, extremely low-latency, scalable networking that uses state-of-the-art technologies such as remote direct-memory access (RDMA), NVIDIA® GPUDirect®, and NVIDIA Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)™. Unlike other networking solutions, NVIDIA Quantum InfiniBand provides self-healing network capabilities, enhanced quality of service (QoS), congestion control, and adaptive routing to provide the highest overall system throughput.

NVIDIA Quantum InfiniBand also dramatically boosts application performance of complex computations while data moves through the data center network. It can participate in the application's runtime, improving performance while reducing the amount of data that traverses the network.

NVIDIA QUANTUM INFINIBAND BENEFITS

- > World's fastest networking, supporting up to 400 gigabits per second (Gb/s) per port
- > Ultra-low and predictable latency
- > Support for NVIDIA GPUDirect offload technology
- > NVIDIA In-Network Computing acceleration engines
- > Adaptive routing
- > Enhanced quality of service (QoS)
- > Reduced costs through a converged data center network

When Time to Insight Is Critical

LAMMPS (Large-Scale Atomic/Molecular Massively Parallel Simulator) is a classical molecular dynamics code. With a focus on materials modeling, LAMMPS is often used for solid-state materials but is also used with soft-matter simulations, such as biomolecules and polymers, nanoparticle coating structures, wetting and surface properties of complex fluids, and more. It can be used for a few particles up to billions. LAMMPS has been used across investigations of DNA structures for decades and has shown improved scalability in large-scale environments.

Because of the extreme efficiency InfiniBand brings to the processing of complex parallel codes, it's the preferred networking solution for molecular dynamics simulations. Figure 1 demonstrates the scalable performance of NVIDIA Quantum 200Gb/s InfiniBand in the rhodopsin test case.

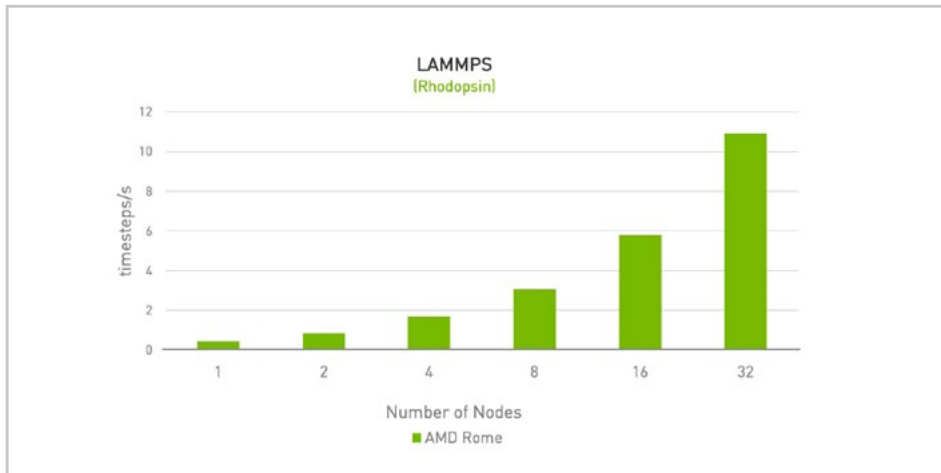


Figure 1. Scalable performance with LAMMPS with NVIDIA Quantum 200Gb/s InfiniBand. System configuration: DELL PowerEdge C6420 server-based 32-node cluster, dual-socket Intel Xeon Gold 6248 CPU @ 2.50GHz, 192GB DDR4 2677MHz RDIMMs per node, NVIDIA ConnectX[®]-6 100Gb/s InfiniBand, NVIDIA Quantum QM8700 switch; OS: RHEL 7.8, NVIDIA MLNX_OFED 4.9, NVIDIA HPC-X 2.7.0, LAMMPS v10-29-2020, Intel Compiler 2020.4.304.

Research Demands Every Ounce of Performance

For single-particle analysis of electron cryogenic microscopy (cryo-EM) data, researchers depend upon RELION, an important tool in the study of living cells. NVIDIA Quantum 200Gb/s InfiniBand achieves superior scalable performance for applications such as RELION. And by taking advantage of NVIDIA SHARP capabilities, as shown in figure 2, researchers can maximize their simulations per day, even when running with a fraction of supercomputing resources.

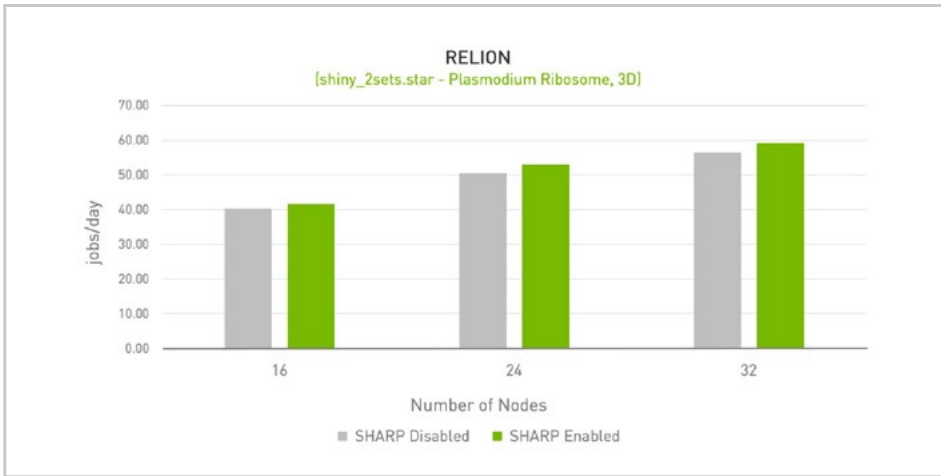


Figure 2. Scalable performance with RELION showcasing benefits of NVIDIA SHARP. System configuration: Supermicro SYS-6029U-TR4 / Foxconn Groot 1A42USF00-600-G 32-node cluster, dual-socket Intel Xeon Gold 6138 CPU @ 2.00GHz, 192GB DDR4 2677MHz RDIMMs per node, NVIDIA ConnectX-6 200Gb/s InfiniBand, NVIDIA Quantum GM8700 switch; OS: RHEL 8.3, NVIDIA MLNX_OFED 5.2.1, NVIDIA HPC-X 2.7.0, Relion 3.1.0.

Faster Networking, Faster Scientific Breakthroughs

For the demanding requirements of bioscience research, NVIDIA delivers unmatched reliability, scalability, and network utilization, allowing researchers to accelerate their research and discovery.

[Learn more](#)

Learn more about how NVIDIA is helping to accelerate scientific discovery at www.nvidia.com/hpc