

NAMD Performance Benchmark and Profiling

October 2009

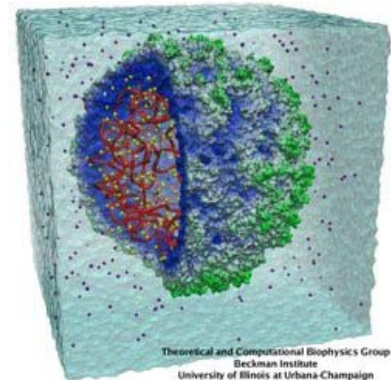
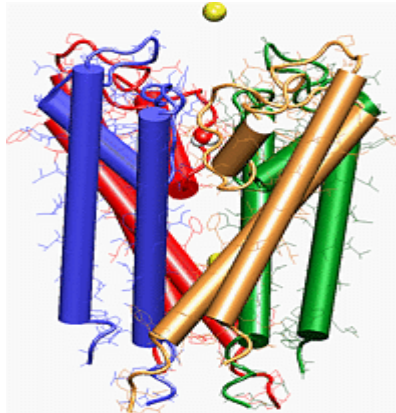


- **This research was performed under the HPC Advisory Council activities**
 - Intel, Dell, Mellanox
 - HPC Advisory Council HPC Center
- **For more info on the listed companies offering, please refer to**
 - www.mellanox.com, www.dell.com/hpc, www.Intel.com
- **For more info on the HPC Advisory Council HPC center, please refer to**
 - http://www.hpcadvisorycouncil.com/cluster_center.php

- A parallel molecular dynamics code that received the 2002 Gordon Bell Award
- Designed for high-performance simulation of large biomolecular systems
 - **Scales to hundreds of processors and millions of atoms**
- Developed by the joint collaboration of the Theoretical and Computational Biophysics Group (TCB) and the Parallel Programming Laboratory (PPL) at the University of Illinois at Urbana-Champaign
- NAMD is distributed free of charge with source code



Theoretical and Computational Biophysics Group
Beckman Institute
University of Illinois at Urbana-Champaign



Theoretical and Computational Biophysics Group
Beckman Institute
University of Illinois at Urbana-Champaign

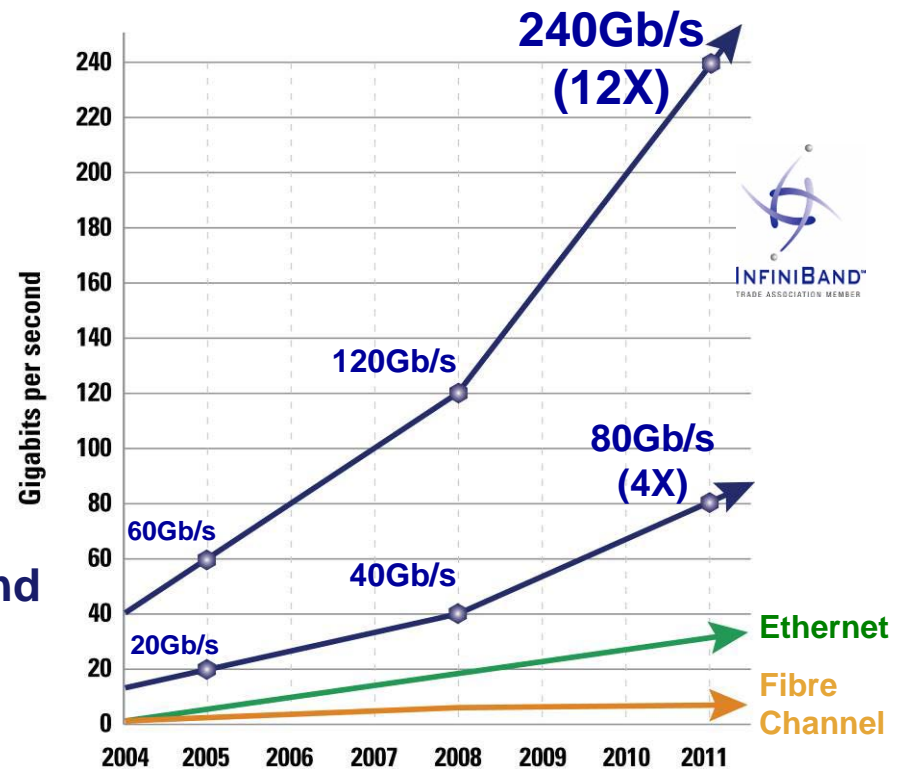
- **This study was done in order to determine:**
 - NAMD performance characteristics
 - NAMD scalability
 - The effect of cluster interconnects on NAMD performance

- **System Structure and Sizing Guidelines**
 - 16-node cluster built with Dell PowerEdge™ M610 blade servers
 - Servers optimized for High Performance Computing environments
 - Building Block Foundations for best price/performance and performance/watt
- **Dell HPC Solutions**
 - Scalable Architectures for High Performance and Productivity
 - Dell's comprehensive HPC services help manage the lifecycle requirements.
 - Integrated, Tested and Validated Architectures
- **Workload Modeling**
 - Optimized System Size, Configuration and Workloads
 - Test-bed Benchmarks
 - ISV Applications Characterization
 - Best Practices & Usage Analysis



- **Industry Standard**
 - Hardware, software, cabling, management
 - Design for clustering and storage interconnect
- **Performance**
 - 40Gb/s node-to-node
 - 120Gb/s switch-to-switch
 - 1us application latency
 - Most aggressive roadmap in the industry
- **Reliable with congestion management**
- **Efficient**
 - RDMA and Transport Offload
 - Kernel bypass
 - CPU focuses on application processing
- **Scalable for Petascale computing & beyond**
- **End-to-end quality of service**
- **Virtualization acceleration**
- **I/O consolidation including storage**

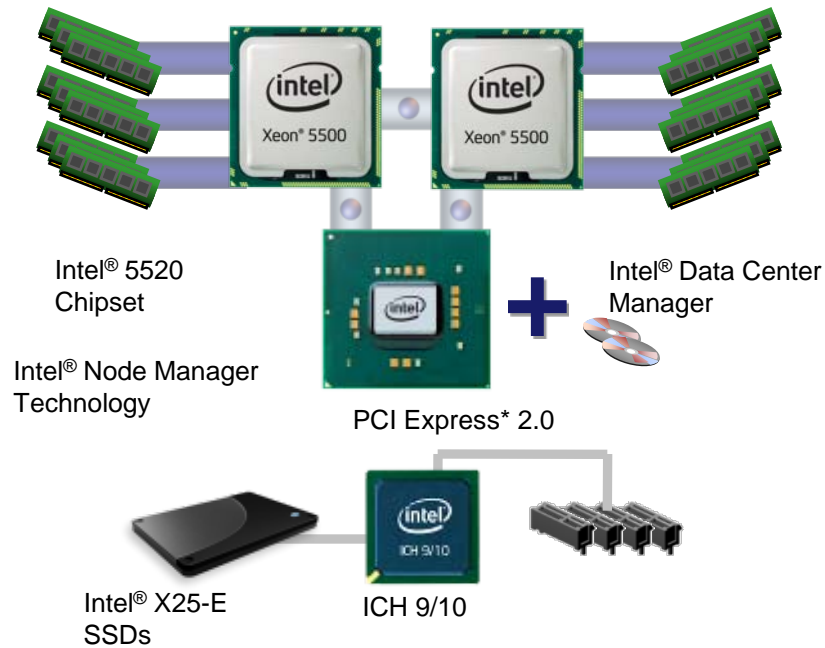
The InfiniBand Performance Gap is Increasing



InfiniBand Delivers the Lowest Latency

Delivering Intelligent Performance

Next Generation Intel® Microarchitecture



Bandwidth Intensive

- Intel® QuickPath Technology
- Integrated Memory Controller

Threaded Applications

- 45nm quad-core Intel® Xeon® Processors
- Intel® Hyper-threading Technology

Performance on Demand

- Intel® Turbo Boost Technology
- Intel® Intelligent Power Technology

Performance That Adapts to The Software Environment

Test Cluster Configuration

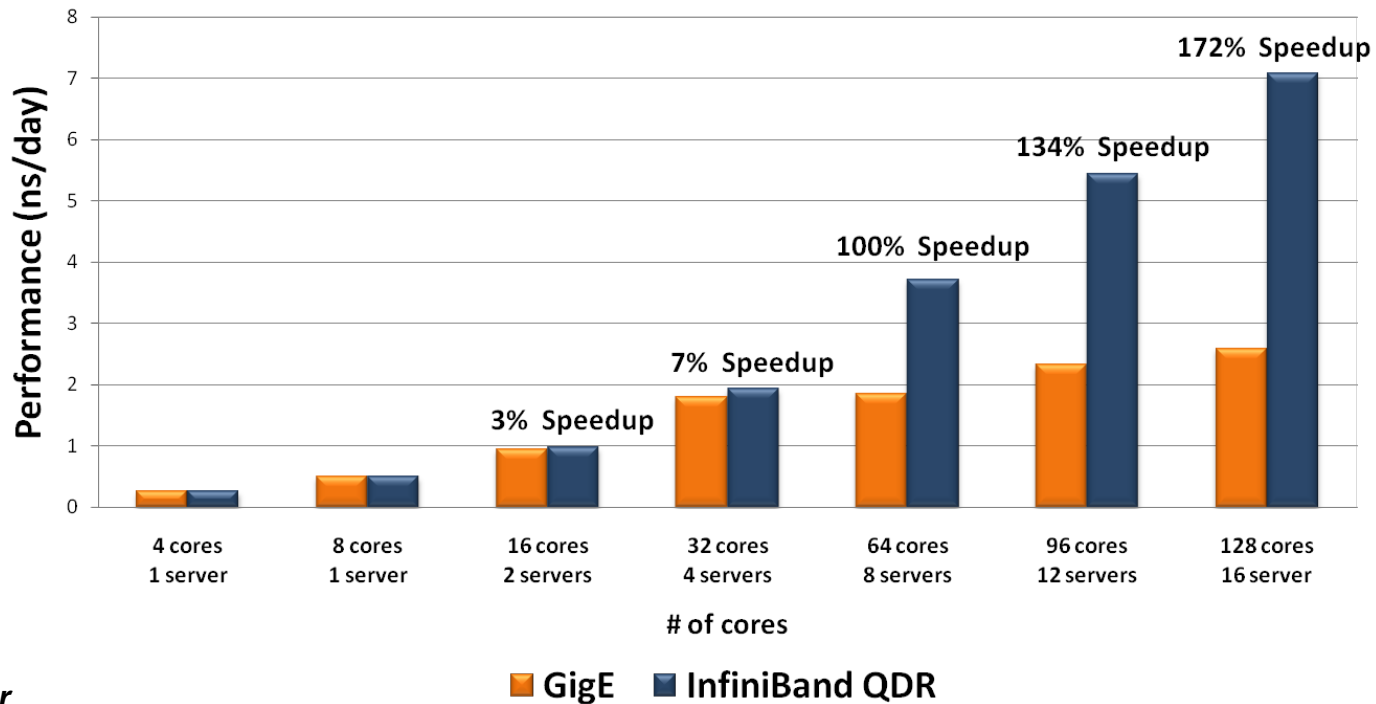
- **Dell™ PowerEdge™ M610 16-node cluster (codename Janus)**
- **Quad-Core Intel X5570 @ 2.93 GHz CPUs**
- **Intel Cluster Ready certified cluster**
- **Mellanox ConnectX MCQH29-XCC 4X QDR InfiniBand mezzanine card**
- **Mellanox M3601Q 32-Port Quad Data Rate (QDR-40Gb) InfiniBand Switch**
- **Memory: 24GB memory per node**
- **OS: RHEL 5.3, OFED 1.4 InfiniBand SW stack**
- **MPI: Open MPI-1.3.3**
- **Application: NAMD version 2.7b1, Charmm-6.1, fftw-2.1.5, gcc-1.4.2**
- **Benchmarks: ApoA1 bloodstream lipoprotein particle model**
 - (92,224 atoms, 12 A cutoff)



NAMD Performance Results

- **The benchmark comprises 92K atoms of lipid, protein, and water**
 - A NAMD data set that models a lipoprotein particle in the bloodstream
- **InfiniBand 40Gb/s outperforms GigE between 16-128 core counts**
 - QDR InfiniBand provides up to 172% higher performance than GigE

NAMD Performance Results



- **Optimal NAMD performance relies on a low latency interconnect**
 - NAMD doesn't scale well on GigE past 32 cores
 - NAMD scales very well on QDR IB up to 128 cores
 - No limitation was found for scaling beyond the given system size
- **InfiniBand enables NAMD performance scalability**
 - InfiniBand QDR performance is up to 172% greater than GigE

Thank You

HPC Advisory Council



All trademarks are property of their respective owners. All information is provided "As-Is" without any kind of warranty. The HPC Advisory Council makes no representation to the accuracy and completeness of the information contained herein. HPC Advisory Council Mellanox undertakes no duty and assumes no obligation to update or correct any information presented herein