



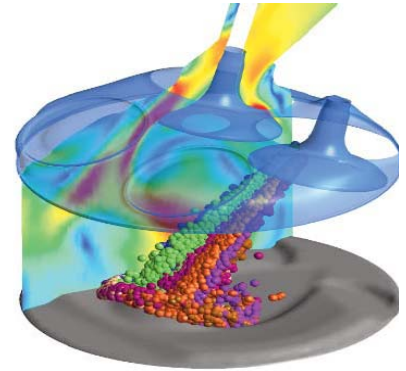
ANSYS FLUENT Performance Benchmark and Profiling

October 2009



- **The following research was performed under the HPC Advisory Council activities**
 - Participating vendors: Intel, ANSYS, Dell, Mellanox
 - Compute resource - HPC Advisory Council Cluster Center
- **The participating members would like to thank ANSYS for their support and guidelines**
- **For more info please refer to**
 - www.mellanox.com, www.dell.com/hpc, www.intel.com,
www.ansys.com

- **Computational Fluid Dynamics (CFD) is a computational technology**
 - Enables the study of the dynamics of things that flow
 - By generating numerical solutions to a system of partial differential equations which describe fluid flow
 - Enables better understanding of qualitative and quantitative physical flow phenomena which is used to improve engineering design
- **CFD brings together a number of different disciplines**
 - Fluid dynamics, mathematical theory of partial differential systems, computational geometry, numerical analysis, computer science
- **ANSYS FLUENT is a leading CFD application from ANSYS**
 - Widely used in almost every industry sector and manufactured product



- **The presented research was done to provide best practices**
 - ANSYS FLUENT performance benchmarking
 - Interconnect performance comparisons
 - Power-efficient simulations
 - Understanding FLUENT communication patterns


Test Cluster Configuration



- **Dell™ PowerEdge™ M610 24-node cluster**
- **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: RHEL5U3, OFED 1.4 InfiniBand SW stack**
- **MPI: HP-MPI 2.3**
- **Application: FLUENT 12.0**
- **Benchmark Workload**
 - **New FLUENT Benchmark Suite**

- **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

- **System Structure and Sizing Guidelines**

- 24-node cluster build with Dell PowerEdge™ M610 blades server
- 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**

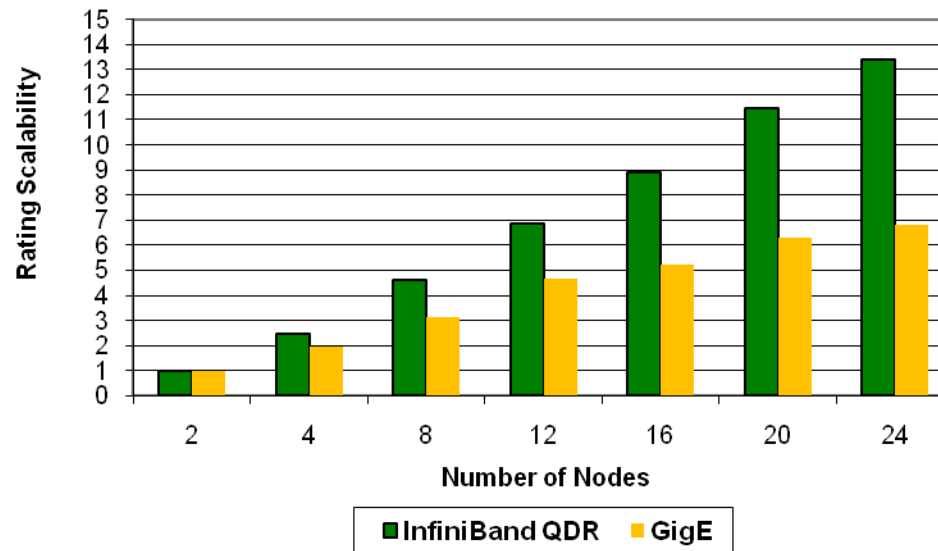
- Optimized System Size, Configuration and Workloads
- Test-bed Benchmarks
- ISV Applications Characterization
- Best Practices & Usage Analysis





- **Input Dataset**
 - Truck_14M (14 millions elements)
 - External Flow Over a Truck Body
- **Performance testing is reported as normalized scalability**
 - The performance multiple increase versus 2 servers as the base line
 - For example, rating scalability of 4 means 4x performance versus the 2 servers case performance
- **InfiniBand QDR delivers above linear scalability**
- **At 24 nodes, GigE wastes 50% of the cluster resources**

Fluent 12.0 Benchmark Result
(Truck_14M)

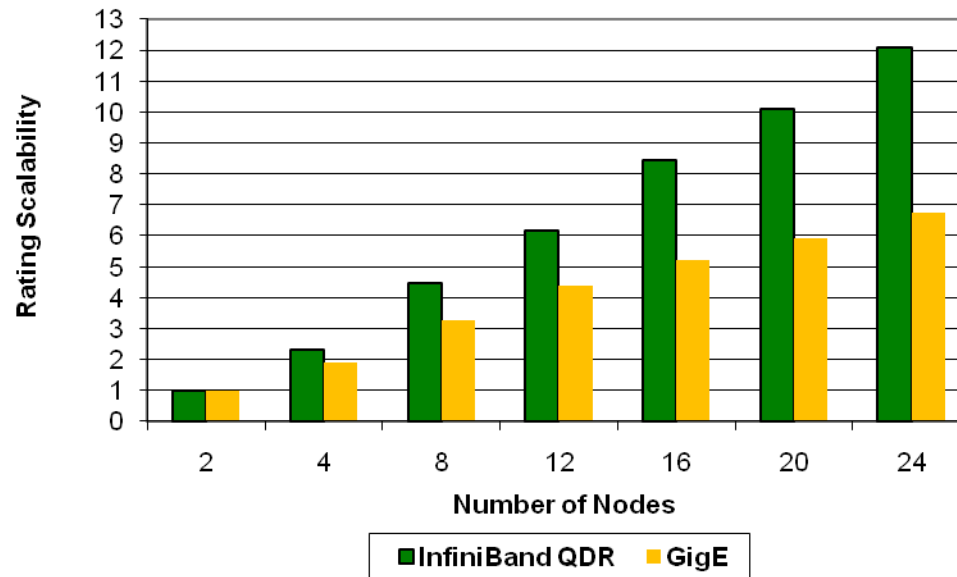


Higher is better



- **Input Dataset**
 - Truck_Poly_14M (14 millions elements)
- **Performance testing is reported as normalized scalability**
 - The performance multiple increase versus 2 servers as the base line
 - For example, rating scalability of 4 means 4x performance versus the 2 servers case performance
- **InfiniBand QDR delivers above linear scalability**
- **At 24 nodes, GigE wastes more than 50% of the cluster resources**

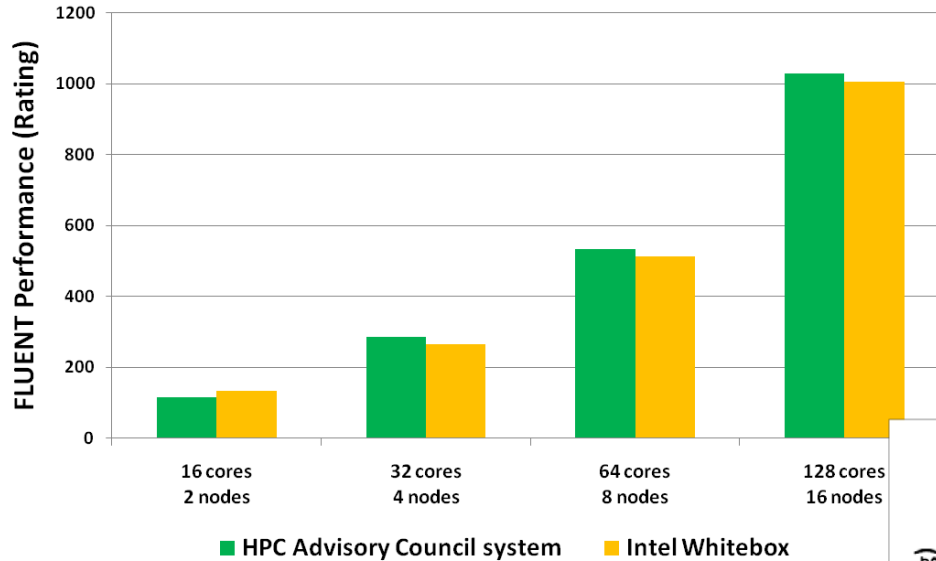
Fluent 12.0 Benchmark Result
(Truck_Poly_14M)



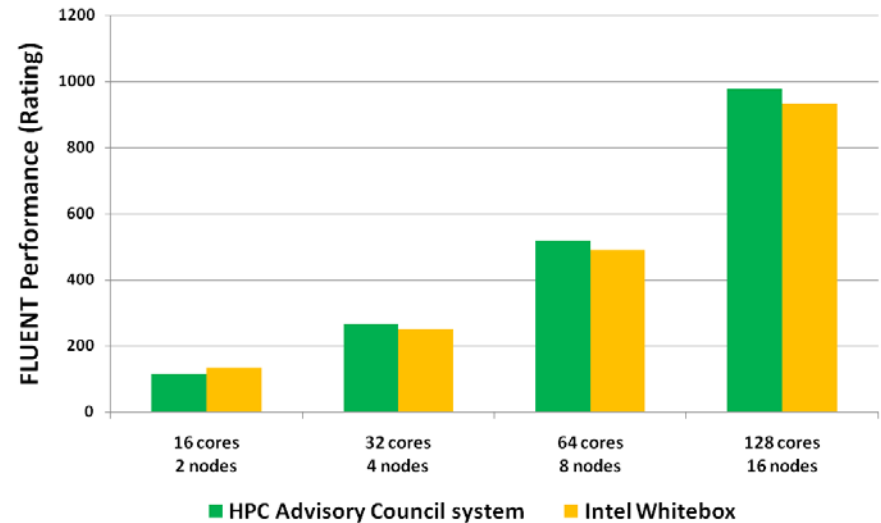
Higher is better

FLUENT Performance Comparison

ANSYS FLUENT Truck_14M Performance



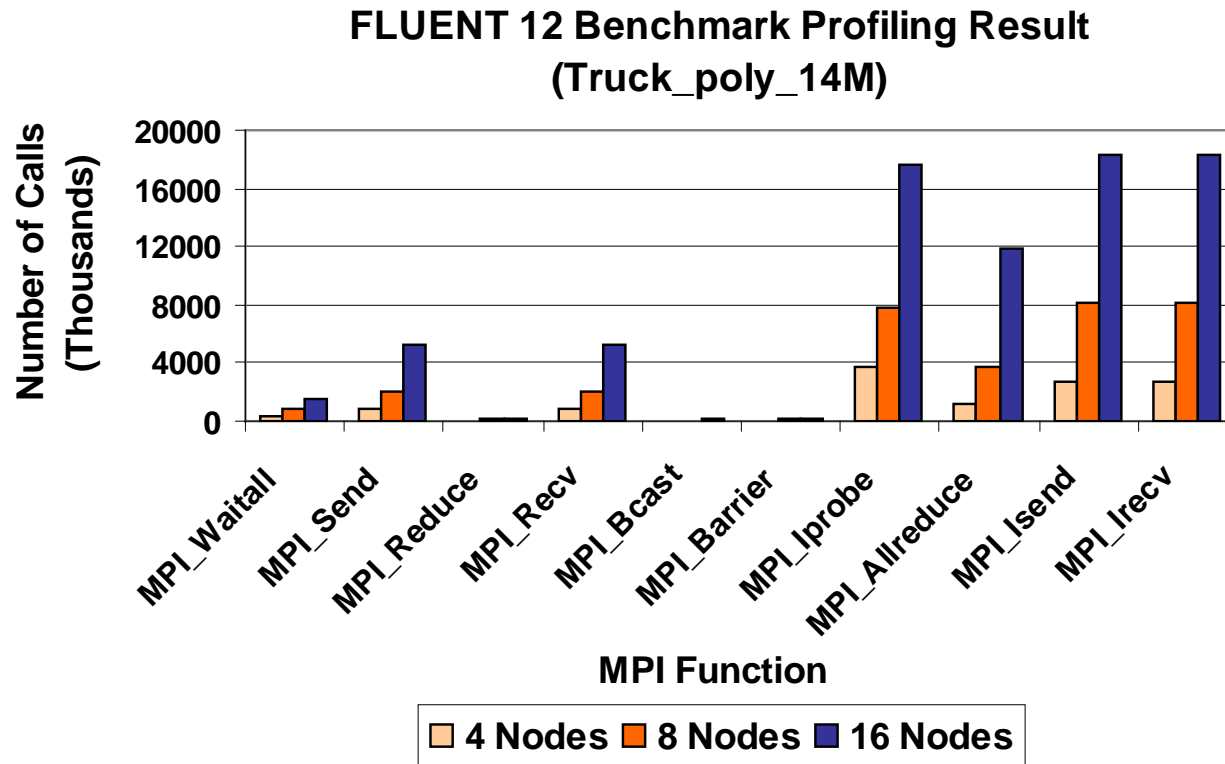
ANSYS FLUENT Truck_Poly_14M Performance



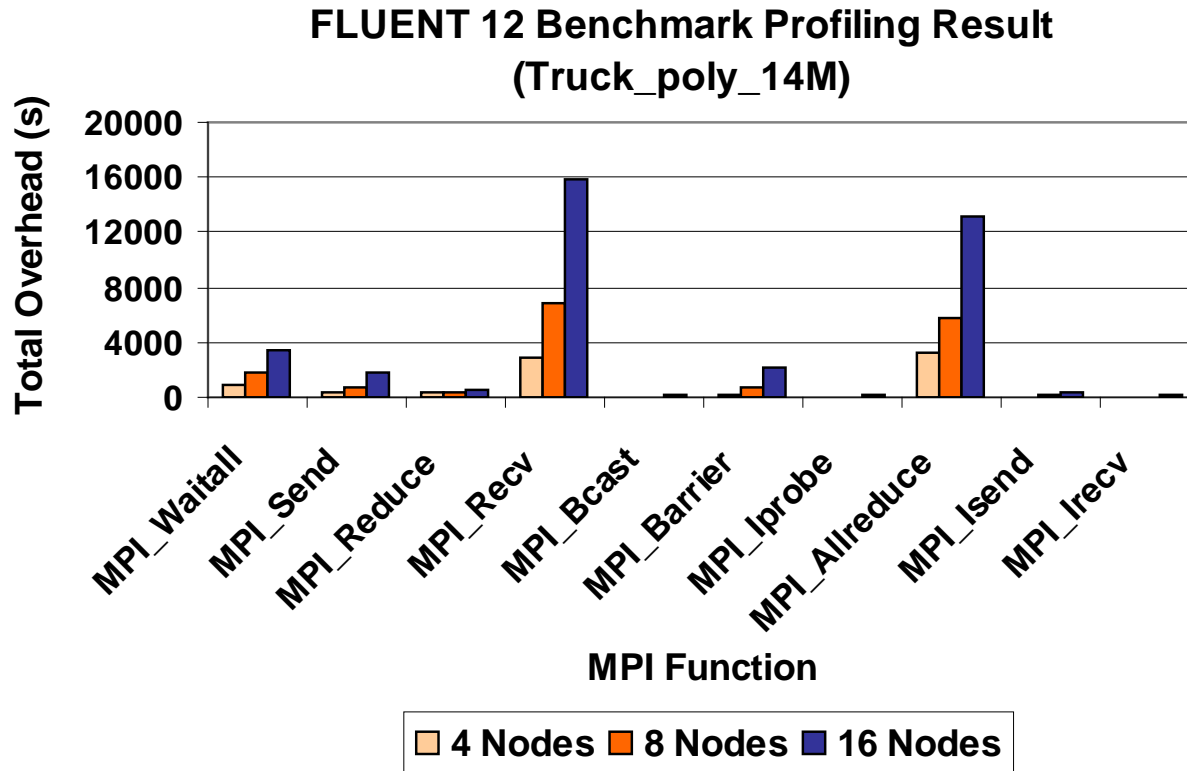
- Performance comparison between the tested results and best know results (Intel white box)
- Comparison demonstrated the capability of balanced system to provide better performance and higher scalability

- **Mostly used MPI functions**

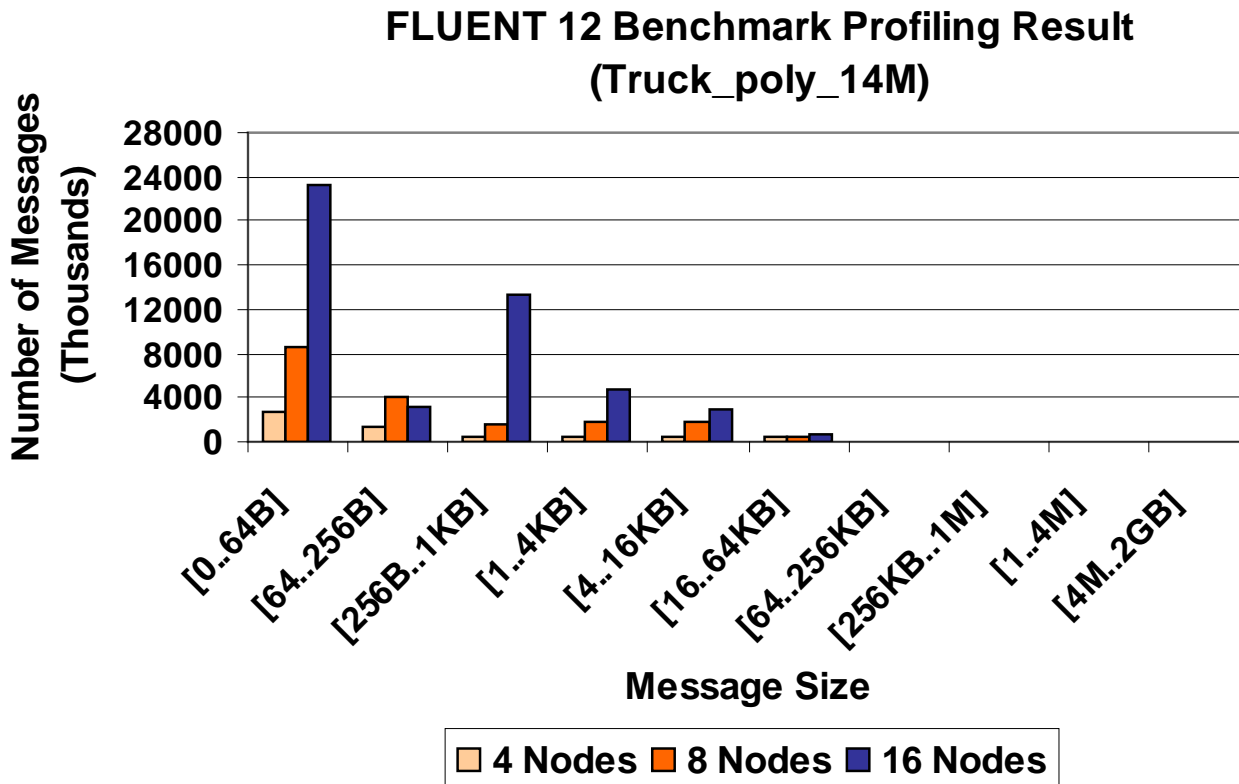
- MPI_Iprobe, MPI_Allreduce, MPI_Isend, and MPI_Irecv



- **MPI_Recv and MPI_Allreduce show highest communication overhead**



- Most data related MPI messages are within 256B-1KB in size
- Typical MPI synchronization messages are lower than 64B in size
- Number of messages increases with cluster size




- **FLUENT 12 was profiled to identify its communication patterns**
- **Frequently used message sizes**
 - 256-1KB messages for data related communications
 - <64B for synchronizations
 - Number of messages increases with cluster size
- **MPI Functions**
 - FLUENT 12 introduced MPI collective functions
 - MPI_Allreduce help improves the communication efficiency
- **Interconnects effect to FLUENT performance**
 - Both interconnect latency (MPI_Allreduce) and throughput (MPI_Recv) highly influence FLUENT performance
 - Further optimization can be made to take bigger advantage of high-speed networks

Productive System = Balanced System

- **Balanced system enables highest productivity**
 - Interconnect performance to match CPU capabilities
 - CPU capabilities to drive the interconnect capability
 - Memory bandwidth to match CPU performance
- **Applications scalability relies on balanced configuration**
 - “Bottleneck free”
 - Each system component can reach its highest capability
- **Dell M610 system integrates balanced components**
 - Intel “Nehalem” CPUs and Mellanox InfiniBand QDR
 - Latency to memory and interconnect latency at the same order of magnitude
 - Provide the leading productivity and power/performance system for FLUENTsimulations

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