

STAR-CCM+ Performance Benchmarking and Profiling

July 2010





- **The following research was performed under the HPC Advisory Council activities**
 - Participating vendors: CD-adapco, HP, Mellanox
 - Compute resource - HPC Advisory Council Cluster Center


- **For more info please refer to**
 - www.mellanox.com
 - <http://www.hp.com/go/hpc>
 - <http://www.cd-adapco.com>

- **STAR-CCM+**

- An engineering process-oriented CFD tool
- Client-server architecture, object-oriented programming
- Delivers the entire CFD process in a single integrated software environment



- **Developed by CD-adapco**




- **The presented research was done to provide best practices**
 - CD-adapco performance benchmarking
 - Interconnect performance comparisons
 - Ways to increase CD-adapco productivity
 - Power-efficient simulations
- **The presented results will demonstrate**
 - The scalability of the compute environment
 - Considerations for power saving through balanced system configuration

- **HP ProLiant SL2x170z G6 16-node cluster**
 - Six-Core Intel X5670 @ 2.93 GHz CPUs
 - Memory: 24GB per node
 - OS: CentOS5U4, OFED 1.5.1 InfiniBand SW stack
- **Mellanox ConnectX-2 adapters and switches**
- **Fulcrum based 10GigE switch**
- **MPI: HP-MPI 2.3.1**
- **Application: STAR-CCM+ Version 5.04.006, STAR-CD Version 4.12.022**
- **Benchmark Workload**
 - STAR-CCM+: Lemans_Poly_17M (Epsilon Euskadi Le Mans car external aerodynamics)

About HP ProLiant SL6000 Scalable System


- **Solution-optimized for extreme scale out**




ProLiant z6000 chassis
Shared infrastructure
– fans, chassis, power



ProLiant SL160z G6 ProLiant SL165z G7
Large memory
-memory-cache apps



ProLiant SL170z G6
Large storage
-Web search and database apps



ProLiant SL2x170z G6
Highly dense
- HPC compute and
web front-end apps

Save on cost and energy -- per node, rack and data center

Mix and match configurations

Deploy with confidence



#1
Power
Efficiency*

* SPECpower_ssj2008
www.spec.org
17 June 2010, 13:28

- **Input Dataset**


- Lemans_Poly_17M

- **InfiniBand QDR enables higher scalability**

- InfiniBand reduces electrical energy/job by 57% or more compared to GigE
- 130% higher performance than GigE at 14 nodes and 66% than 10GigE at 8 nodes



STAR-CCM+ Benchmark (Lemans_Poly_17M)



Higher is better

12 Cores/Node

- **Input Dataset**


- Lemans_Poly_17M

- **Core/socket comparison**

- Using 6 cores/socket (12 cores/node) provides higher performance at all node count
- Up to 21% higher performance than 8 cores/node



STAR-CCM+ Benchmark
(Lemans_Poly_17M)



Higher is better


InfiniBand

- **Interconnect comparison shows**
 - InfiniBand delivers superior performance in every cluster size
 - Low latency InfiniBand enables better scalability than 10GigE and GigE
- **Cores/Socket Usage**
 - Fully utilizing all the cores in each node yields the best performance
- **InfiniBand QDR saves power**
 - STAR-CCM+
 - InfiniBand reduces power consumption/job by 57% or more compared to GigE


• Runtime Distribution

- MPI communication overhead remains small with InfiniBand as cluster size scales
- MPI overhead becomes dominated with GigE as node count increases
- MPI overhead over 10GigE grows with cluster size


Runtime Distribution (InfiniBand QDR)



Runtime Distribution (GigE)



Runtime Distribution (10GigE)




12 Cores/Node


- **Major MPI Functions**

- MPI_Testall, MPI_Allreduce, MPI_Bcast, and MPI_Reduce
- MPI_Testall is testing the completion of no-blocking send/recv
- MPI_Bcast (Collectives) percentage increase as cluster size scales


96 Cores



144 Cores



- Majority messages are small messages
- Number of messages increases as cluster size scales



- **STAR-CCM+ was profiled to understand its communication pattern**
- **Most used message sizes**
 - Small messages
 - Total number of messages increases with cluster size
- **Interconnects effect to STAR-CCM+ performance**
 - Interconnect latency is critical to application performance
 - Communication overhead remains small with InfiniBand
 - GigE creates big communication overhead for application

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