


# Quantum ESPRESSO Performance Benchmarking and Profiling

Sept 2010



- **The following research was performed under the HPC Advisory Council activities**
  - Participating vendors: HP, Mellanox
  - Compute resource - HPC Advisory Council Cluster Center
  
- **For more info please refer to**
  - [www.mellanox.com](http://www.mellanox.com)
  - <http://www.hp.com/go/hpc>
  - <http://www.quantum-espresso.org>

- Quantum ESPRESSO stands for opEn Source Package for Research in Electronic Structure, Simulation, and Optimization
- It is an integrated suite of computer codes for electronic-structure calculations and materials modeling at the nanoscale
- It is based on
  - Density-functional theory
  - Plane waves
  - Pseudopotentials (both norm-conserving and ultrasoft)
- Open source under the terms of the GNU General Public License




- **The presented research was done to provide best practices**
  - Quantum ESPRESSO performance benchmarking
    - Interconnect performance comparisons
  - Ways to increase ESPRESSO 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: Open MPI 1.4.1**
- **Application: Quantum ESPRESSO 4.1.2**
- **Benchmark Workload**
  - Medium size DEISA benchmark AUSURF112
    - Gold surface (112 atoms)

# 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

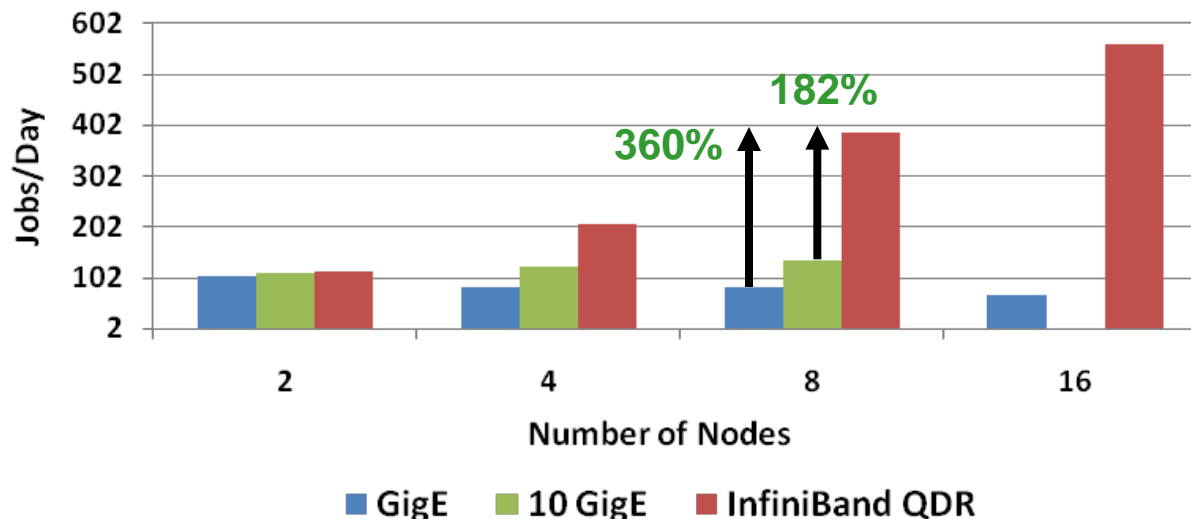


#1  
Power  
Efficiency\*

\* SPECpower\_ssj2008  
[www.spec.org](http://www.spec.org)  
17 June 2010, 13:28

- **Input Dataset**
  - DEISA benchmark AUSURF112
- **InfiniBand QDR enables higher scalability**
  - 360% higher performance than GigE at 8 nodes
  - 182% higher performance than 10GigE at 8 nodes
  - GigE stops scaling after 2 nodes
- **InfiniBand reduces electrical energy/job**
  - by 80% or more compared to GigE and 65% compared to 10GigE

## Quantum ESPRESSO Benchmark (AUSURF112)



*Higher is better*

**12 Cores/Node**


- **Input Dataset**

- DEISA benchmark AUSURF112

- **Performance comparison**

- 1 job mode: All cores per node used by single job
- 2 jobs mode: Each job runs over half number of cores per socket
- Running 2 jobs in parallel delivers 53% higher productivity than single job alone
  - Performance advantage grows as cluster size increases

## Quantum ESPRESSO Benchmark Results (AUSURF112)




*Higher is better*

**192 cores/16Nodes**



# Quantum ESPRESSO MPI Profiling – MPI Time

- **MPI\_Barrier, MPI\_Alltoallv, and MPI\_Addreduce generates most communication overhead**
- **MPI\_Allreduce and MPI\_Alltoallv overhead grow faster than other function**



- **Interconnect comparison shows**
  - InfiniBand delivers superior performance in every cluster size
  - Low latency InfiniBand enables much higher scalability than 10GigE and GigE
- **Customized job placement increases application productivity**
  - 53% more jobs can be completed by running 2 jobs concurrently
- **InfiniBand QDR saves power**
  - Reduces power consumption/job by
    - 80% or more compared to GigE
    - 65% or more compared to 10GigE
- **MPI Profiling shows interconnect latency is the key to enable Quantum ESPRESSO scalability**

# 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