

Abaqus 6.12 Performance Benchmark and Profiling

September 2012



- **The following research was performed under the HPC Advisory Council activities**

- Special thanks for: HP, Mellanox







- **For more information on the supporting vendors solutions please refer to:**

- www.mellanox.com, <http://www.hp.com/go/hpc>

- **For more information on the application:**

- <http://www.simulia.com>

- **Abaqus Unified FEA product suite offers powerful and complete solutions for both routine and sophisticated engineering problems covering a vast spectrum of industrial applications**
- **The Abaqus analysis products listed below focus on:**
 - Nonlinear finite element analysis (FEA)
 - Advanced linear and dynamics application problems
- **Abaqus/Standard**
 - General-purpose FEA that includes broad range of analysis capabilities
- **Abaqus/Explicit**
 - Nonlinear, transient, dynamic analysis of solids and structures using explicit time integration



- **The presented research was done to provide best practices**
 - Abaqus performance benchmarking
 - Interconnect performance comparisons
 - MPI performance comparison
 - Understanding Abaqus communication patterns

- **The presented results will demonstrate**
 - The scalability of the compute environment to provide nearly linear application scalability

- **HP ProLiant SL230s Gen8 4-node “Athena” cluster**
 - Processors: Dual Eight-Core Intel Xeon E5-2680 @ 2.7 GHz
 - Memory: 32GB per node, 1600MHz DDR3 DIMMs
 - OS: RHEL 6 Update 2, OFED 1.5.3 InfiniBand SW stack
- **Mellanox ConnectX-3 VPI InfiniBand adapters**
- **Mellanox SwitchX SX6036 56Gb/s InfiniBand and 40Gb/s Ethernet Switch**
- **MPI: Platform MPI 8.1.2 (vendor provided)**
- **Application: Abaqus 6.12-2**
- **Benchmark Workload:**
 - Abaqus/Explicit benchmarks: E6: Concentric Spheres



About HP ProLiant SL230s Gen8

Item	SL230 Gen8
Processor	Two Intel® Xeon® E5-2600 Series, 4/6/8 Cores,
Chipset	Intel® Sandy Bridge EP Socket-R
Memory	(512 GB), 16 sockets, DDR3 up to 1600MHz, ECC
Max Memory	512 GB
Internal Storage	Two LFF non-hot plug SAS, SATA bays or Four SFF non-hot plug SAS, SATA, SSD bays Two Hot Plug SFF Drives (Option)
Max Internal Storage	8TB
Networking	Dual port 1GbE NIC/ Single 10G NIC
I/O Slots	One PCIe Gen3 x16 LP slot 1Gb and 10Gb Ethernet, IB, and FlexFabric options
Ports	Front: (1) Management, (2) 1GbE, (1) Serial, (1) S.U.V port, (2) PCIe, and Internal Micro SD card & Active Health
Power Supplies	750, 1200W (92% or 94%), high power chassis
Integrated Management	iLO4 hardware-based power capping via SL Advanced Power Manager
Additional Features	Shared Power & Cooling and up to 8 nodes per 4U chassis, single GPU support, Fusion I/O support
Form Factor	16P/8GPUs/4U chassis



- **Input dataset: E6**
 - Concentric Spheres
- **Intel E5-2680 processors (Sandy Bridge) cluster outperforms prior CPU generation**
 - Performs 53% higher than X5670 cluster at 4 nodes
- **System components used:**
 - Athena: 2-socket Intel E5-2680 @ 2.7GHz, 1600MHz DIMMs, FDR InfiniBand, 1HDD
 - Plutus: 2-socket Intel X5670 @ 2.93GHz, 1333MHz DIMMs, QDR InfiniBand, 1HDD

Abaqus/Explicit Performance (E6)




Lower is better

Wall Clock Time:
based on total
runtime

- **InfiniBand FDR is the most efficient inter-node communication for Abaqus/Explicit**
 - Outperforms 1GbE by 109% at 4 nodes
 - Outperforms 10GbE by 23% at 4 nodes
 - Outperforms 40GbE by 20% at 4 nodes

Abaqus/Explicit Performance (E6)




■ 1GigE ■ 10GbE ■ 40GbE ■ InfiniBand FDR

*Performance Rating:
based on total
runtime*


Higher is better

- **InfiniBand FDR reduces the MPI communication time**
 - InfiniBand FDR consumes about 35% of total runtime at 4 nodes
 - Ethernet solutions consume from 48% to 70% at 4 nodes


Abaqus/Explicit Profiling
(S4B, 4 nodes, InfiniBand FDR)
MPI/User Time Ratio




Abaqus/Explicit Profiling
(S4B, 4 nodes, 1GbE)
MPI/User Time Ratio



Abaqus/Explicit Profiling
(S4B, 4 nodes, 10GbE)
MPI/User Time Ratio



Abaqus/Explicit Profiling
(S4B, 4 nodes, 40GbE)
MPI/User Time Ratio

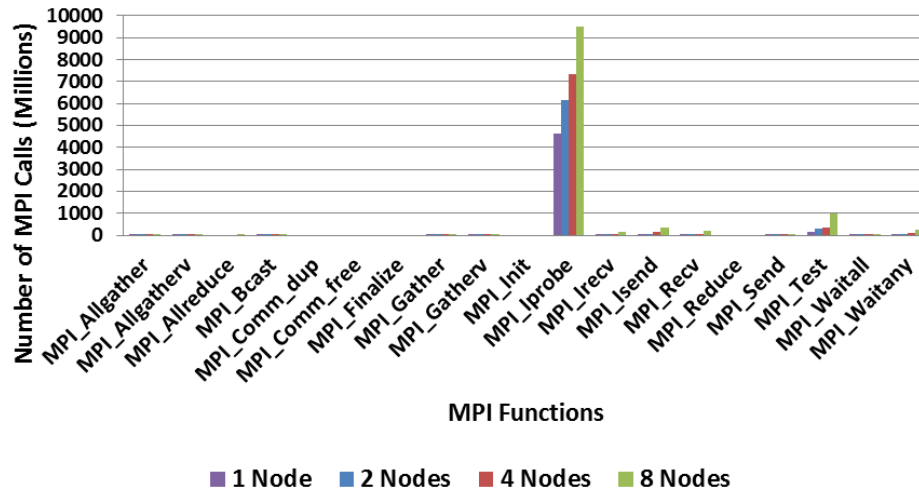


Abaqus/Explicit Profiling – Time Spent in MPI

- **Abaqus/Explicit shows high usage for testing non-blocking messages**
 - MPI_Iprobe (95%), MPI_Test (3%)

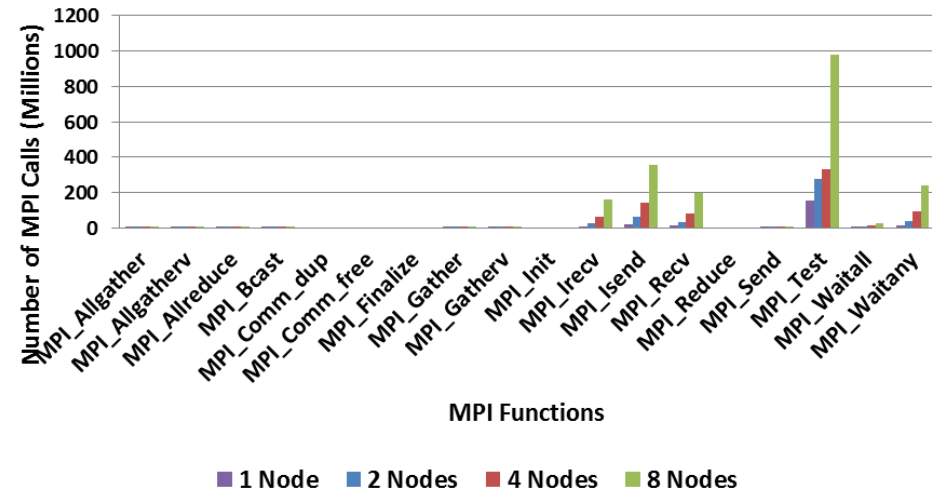
With all MPI calls

Abaqus/Explicit Profiling (E6)
Number of MPI Calls



Without MPI_Iprobe

Abaqus/Explicit Profiling (E6)
Number of MPI Calls




Abaqus/Explicit Profiling – Time Spent in MPI

- **Abaqus/Explicit: More time spent on MPI collective operations:**
 - InfiniBand FDR: MPI_Gather(82%), MPI_Allreduce (4%), MPI_Scatterv(3%)
 - 1GbE: MPI_Gather(41%), MPI_Irecv(15%), MPI_Scatterv(12%)

Abaqus/Explicit Profiling

(E6, 4-node, 1GbE)

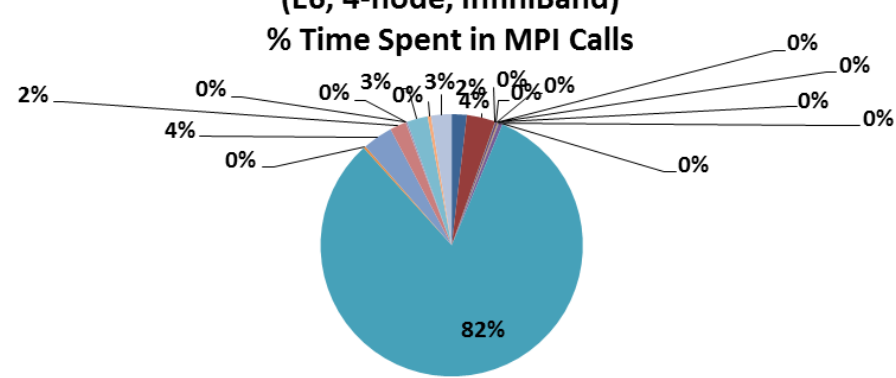
% Time Spent in MPI Calls



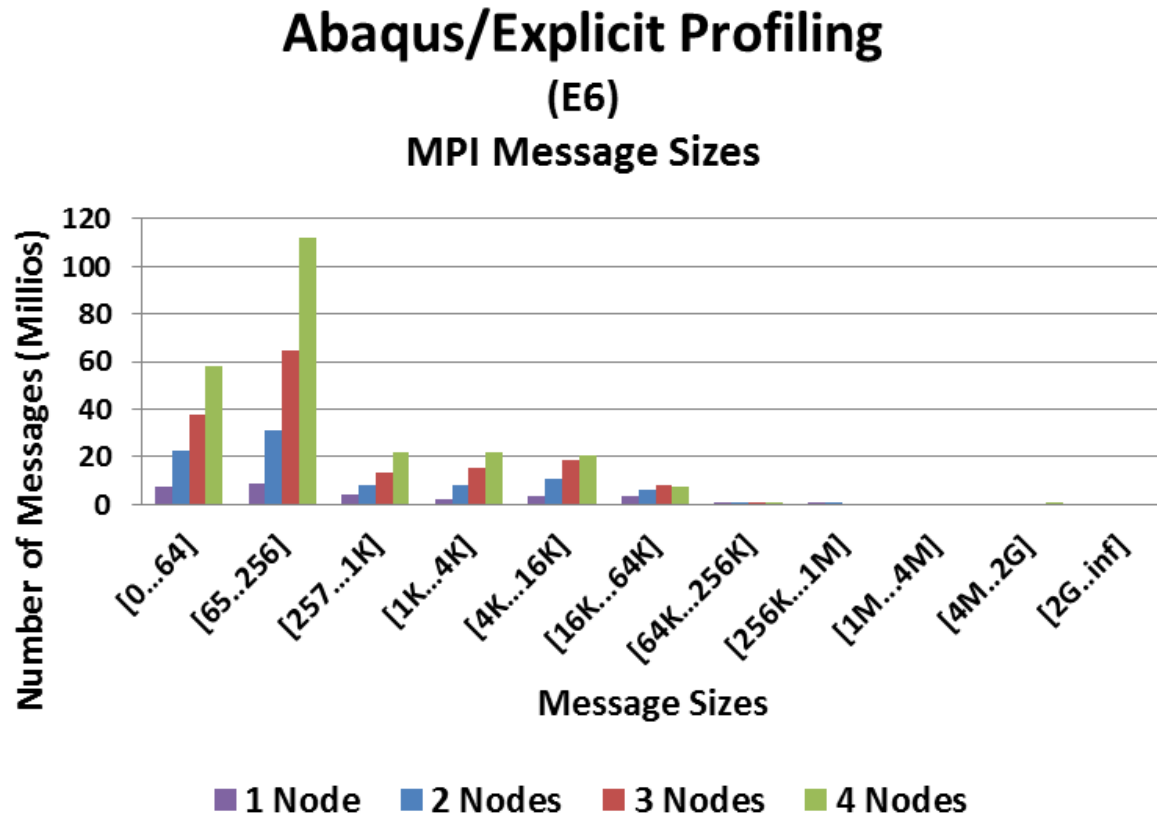
Abaqus/Explicit Profiling

(E6, 4-node, InfiniBand)

% Time Spent in MPI Calls



- **Abaqus/Explicit shows a wide distribution of small message sizes**
 - Small messages peak in the range from 65B to 256B



- **HP ProLiant Gen8 servers delivers better performance than its predecessor**
 - ProLiant Gen8 equipped with Intel E5 series processes and InfiniBand FDR
 - Up to 53% higher performance than ProLiant G7 when compared at 4 nodes
- **InfiniBand FDR is the most efficient inter-node communication for Abaqus/Explicit**
 - Outperforms 1GbE by 109% at 4 nodes
 - Outperforms 10GbE by 23% at 4 nodes
 - Outperforms 40GbE by 20% at 4 nodes
- **Abaqus Profiling**
 - InfiniBand FDR reduces communication time; provides more time for computation
 - InfiniBand FDR consumes 35-40% of total time, versus 45-70% for Ethernet solutions
 - MPI:
 - Large MPI call volumes for testing non-blocking data transfers (MPI_Iprobe, MPI_Test, etc)
 - MPI time is spent mostly on collective operations
 - Messages are concentrated in small messages, peak at around 65-256 bytes

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