HPC ADVISORY COUNCIL

MSC Nastran Performance Benchmark and Profiling

August 2012

hp

Mellanox TECHNOLOGIES

MSC Software

Note

- **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.mscsoftware.com/

MSC Nastran Application

- **MSC Nastran is a widely used Finite Element Analysis (FEA) solver**
- **Used for simulating stress, dynamics, or vibration of real-world, complex systems**
- **Nearly every spacecraft, aircraft, and vehicle designed in the last 40 years has been analyzed using MSC Nastran**

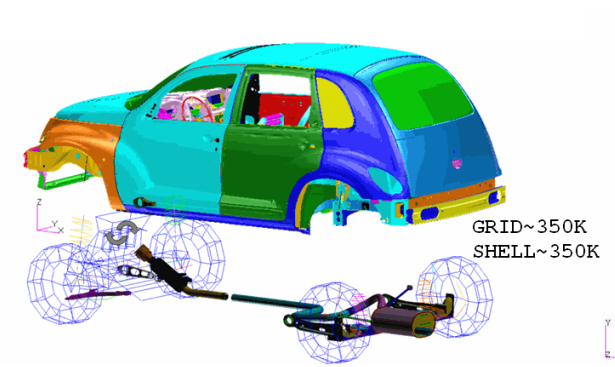

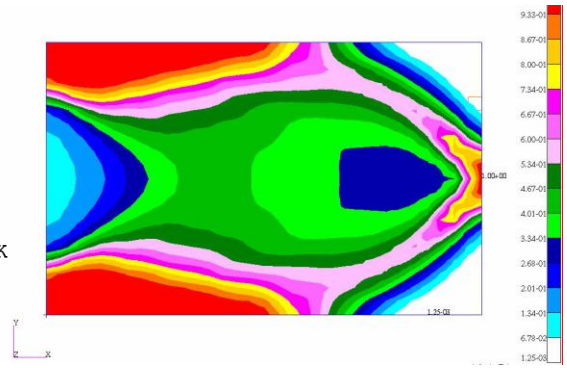

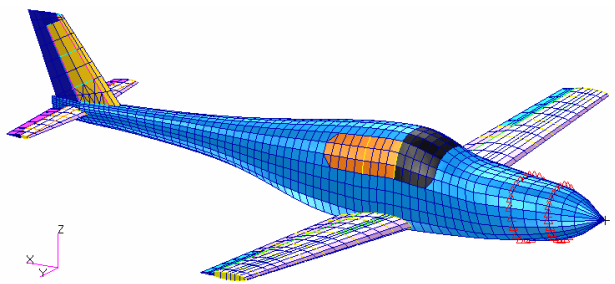

Objectives

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

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

Test Cluster Configuration

- **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, Intel MPI 4.0.3**
- **Application: MSC Nastran 2012.1**
- **Benchmark Workload:**
- **Input dataset:**
  - xl0tdf1: Power Train (Ndof: 529,257, SOL108, Direct Freq)
  - md0mdf1: Half Sphere (Ndof: 42,066 SOL 108 , Direct Freq w/ Umfpack)

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 FlexF abric 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

MSC Nastran Result (xl0tdf1)

- **Input dataset: xl0tdf1**
  - Car Body (Ndof 529,257, SOL111, Direct Frequency Response)
  - Memory: 520MB, SCR Disk: 5GB, Total I/O 190GB
- **Time reduced as more nodes are being utilized for computation**
  - Up to 74% in time saved by running on a 4 DMPs cluster versus 1 DMP

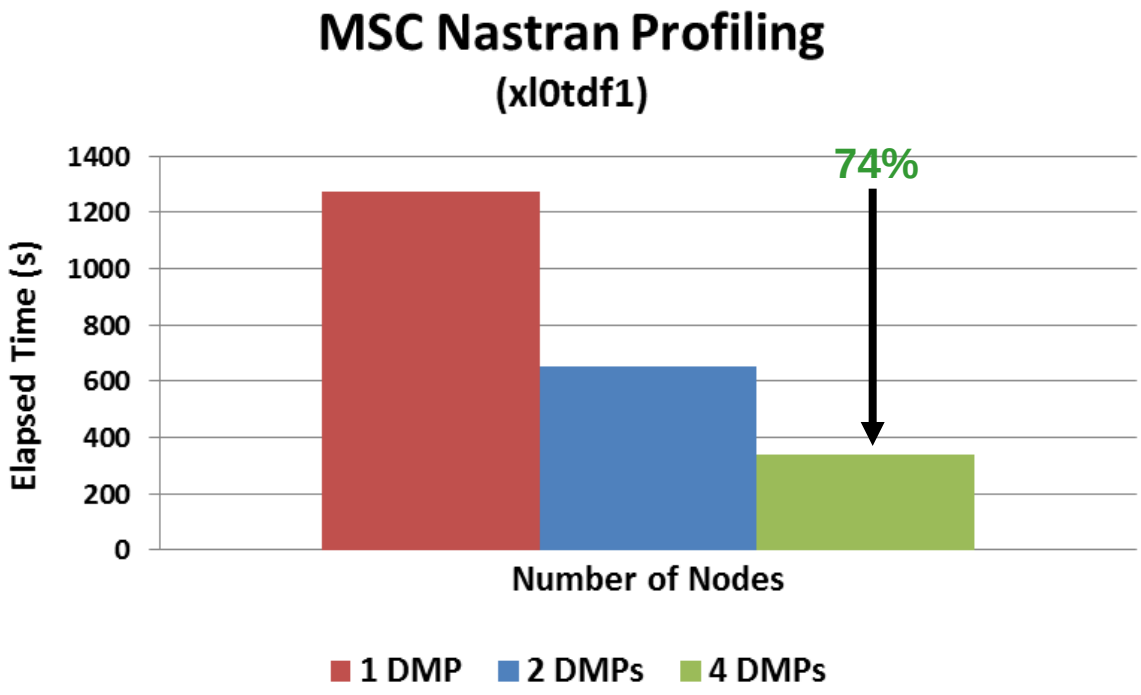

Lower is better

MSC Nastran Result (md0mdf1)

- **Input dataset: md0mdf1**
  - Half Sphere (Ndof 42,066, SOL108, Direct Freq w/ Umfpack)
  - Memory: 1GB, Total I/O 0.1GB
- **Time reduced as more nodes are being utilized for computation**
  - Up to 91% in time saved by running on a 16 DMPs versus 1 DMP

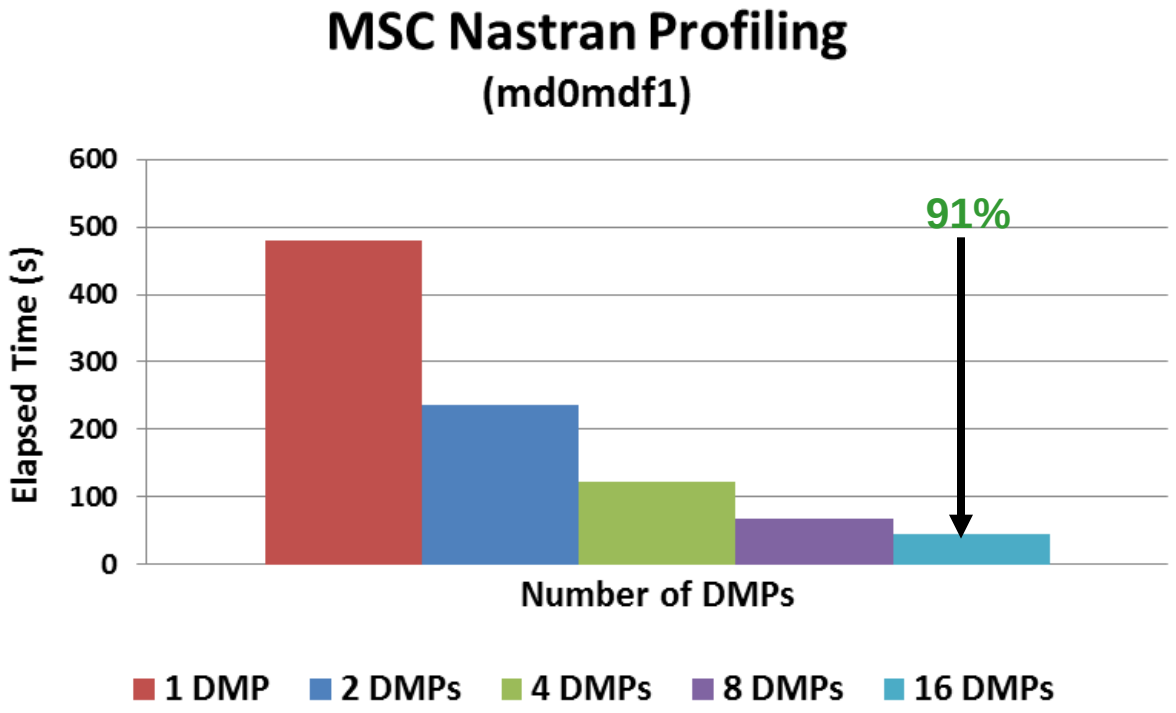

Lower is better

SMP=1

MSC Nastran Result – Comparison (md0dmf1)

- **Intel E5-2600 Series cluster delivers higher performance**
  - Up to 46% in higher performance than best published result* at 8 DMPs
- **Reference hardware configuration:**
  - Dual Intel Westmere X5682 @ 3.47GHz, Linux RHEL 6.1, 96GB 1333MHz DDR3 DIMMs

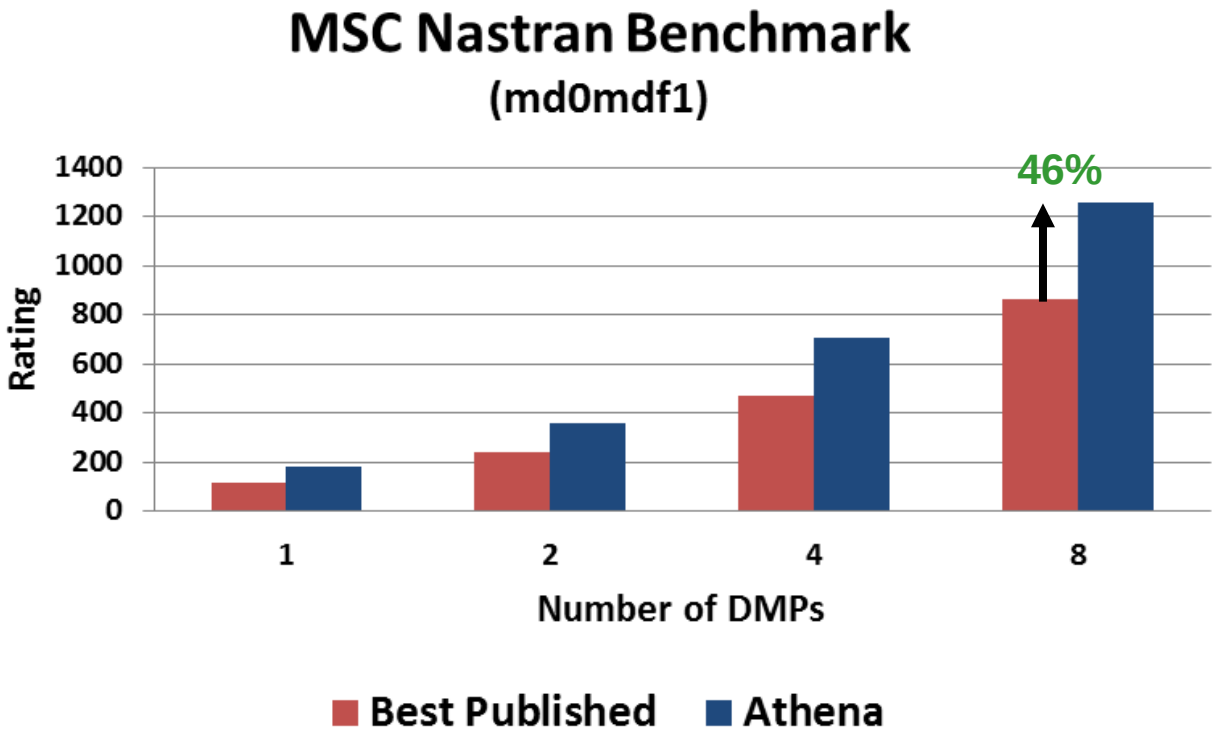

Higher is better *http://www.mscsoftware.com/support/prod_support/nastran/performance/msc20121_par.cfm ***SMP=1***

MSC Nastran Benchmark – Comparision (xl0tdf1)

- **Intel E5-2600 Series cluster delivers higher performance**
  - Up to 39% in higher performance than best published result* at 4 DMPs
- **Reference hardware configuration:**
  - Dual Intel Westmere X5682 @ 3.47GHz, Linux RHEL 6.1, 96GB 1333MHz DDR3 DIMMs

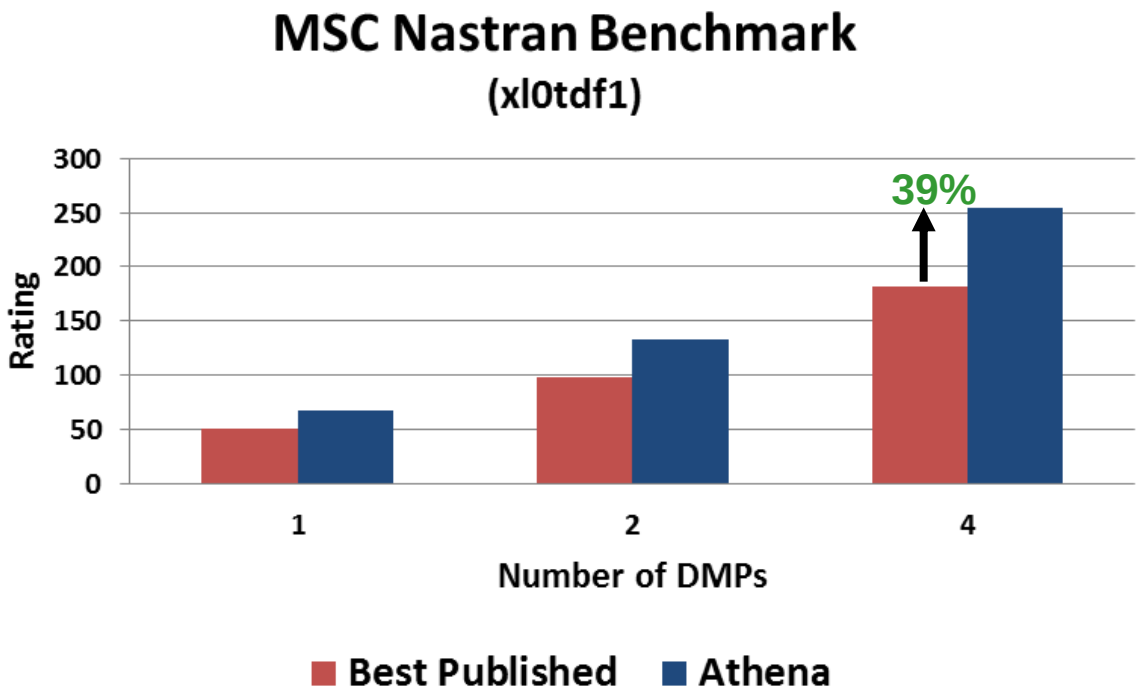

Higher is better

SMP=1

MSC Nastran Benchmark – MPI

- **Both Platform MPI and Intel MPI performs about the same**
  - Intel MPI delivers slightly better performance

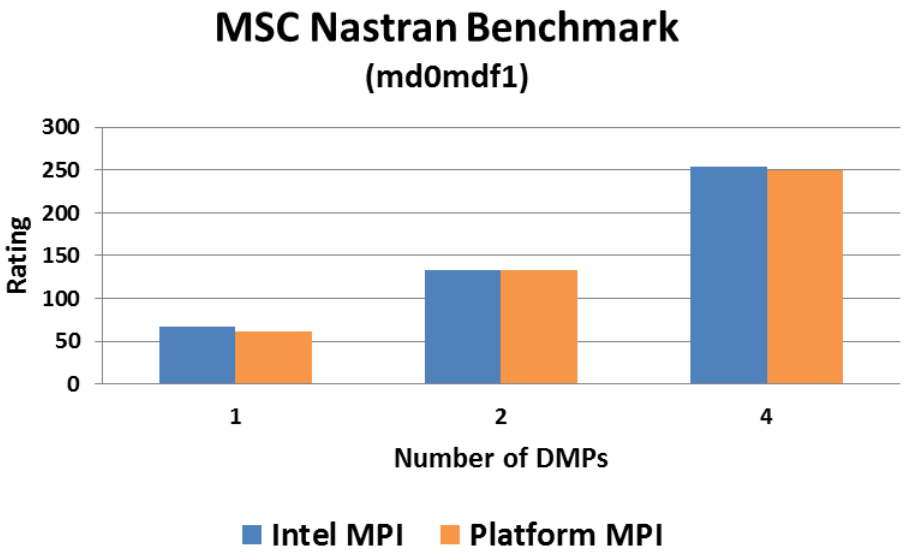

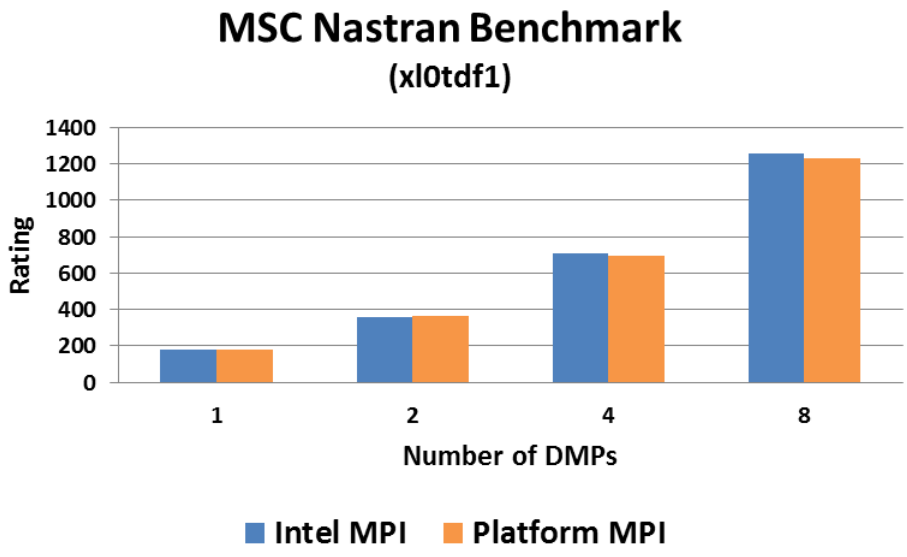

Higher is better

InfiniBand FDR

MSC Nastran Profiling – MPI Functions

- **Mostly used MPI functions**
  - MPI_Recv (50%) and MPI_Ssend (50%)
- **Mostly time consuming MPI functions**
  - MPI_Recv (61%) follows by MPI_Ssend (15%)

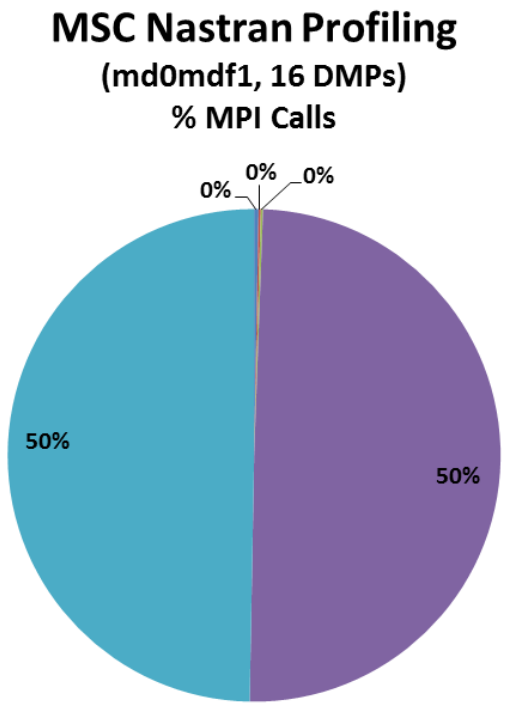

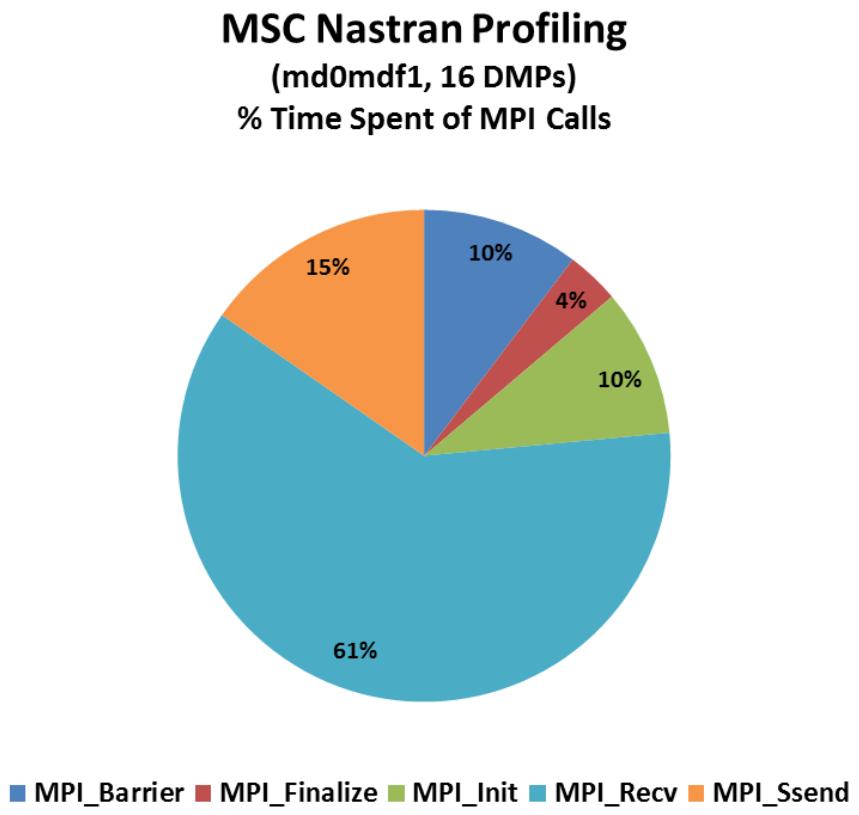

MSC Nastran Profiling – Message Size

- **Majority of MPI messages are small messages**
  - In the range of 0-64 bytes

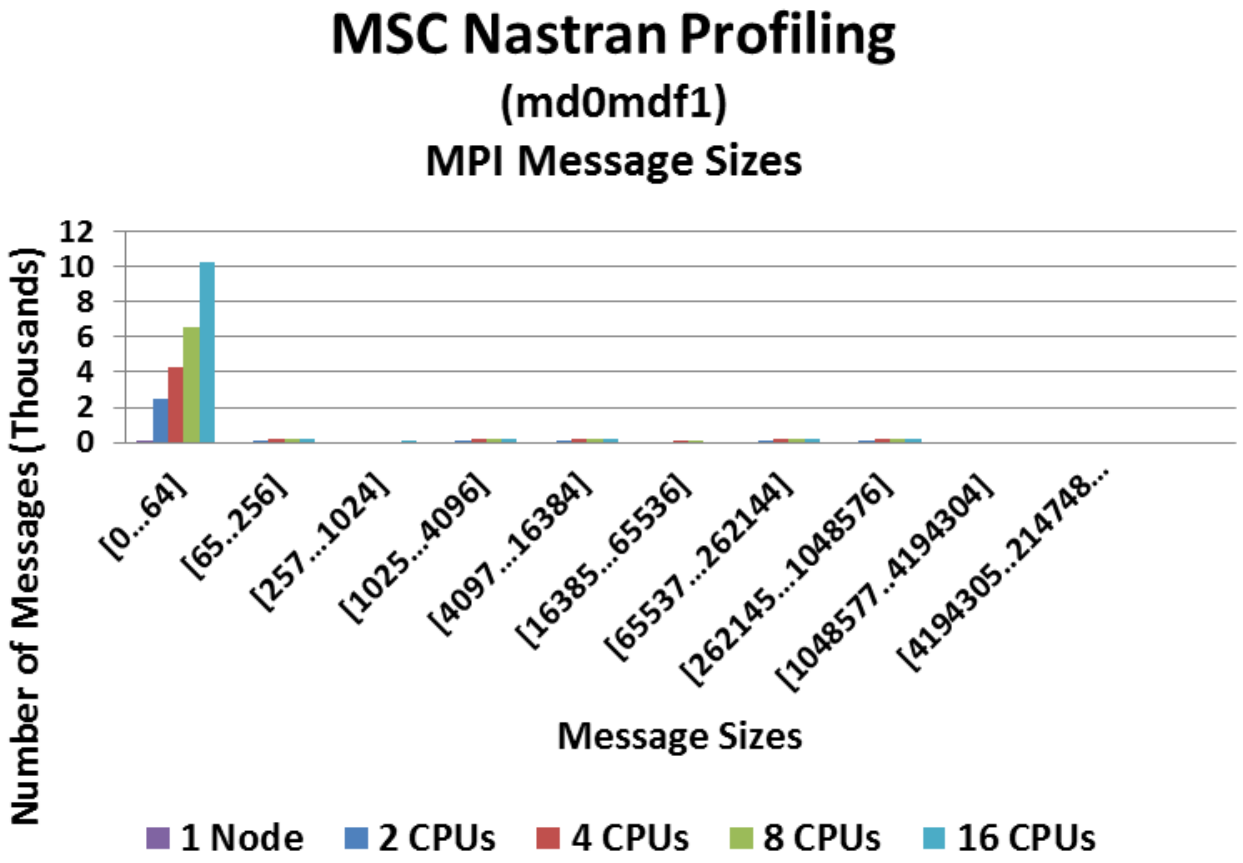

MSC Nastran Profiling – Disk IO

- **Heavy disk write access is seen throughout the test run**
  - Not much access for disk IO reads
  - Tests shows that Nastran could benefit from better disk IO

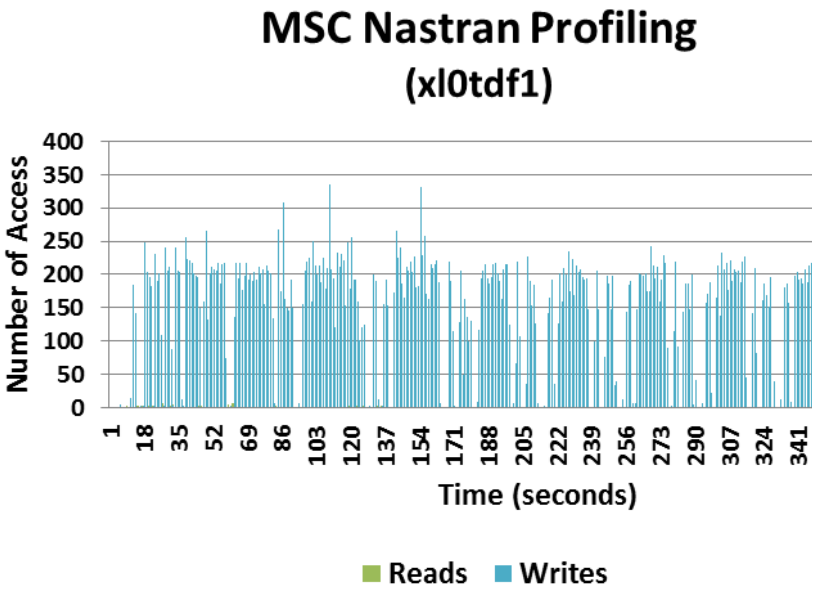

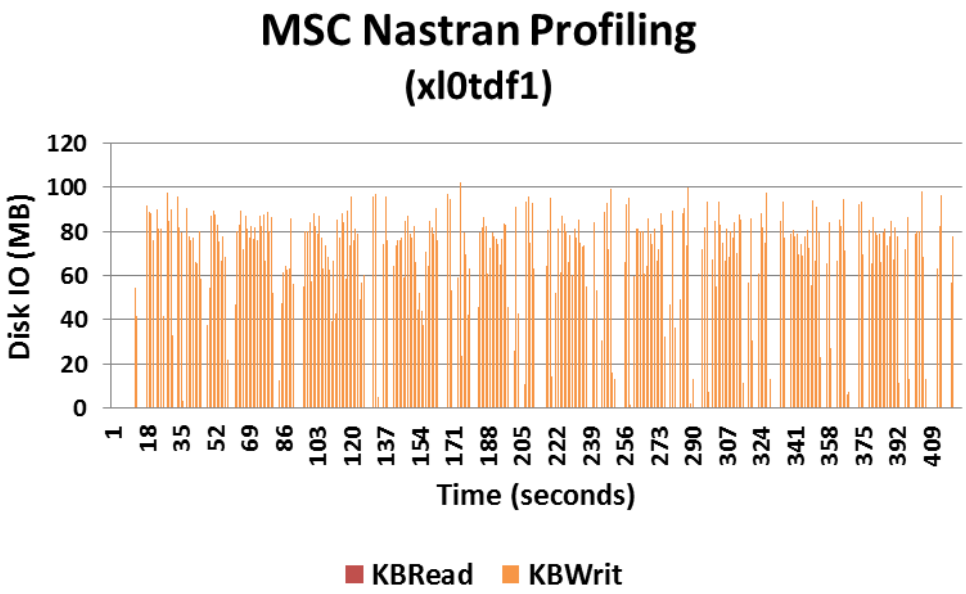

MSC Nastran Summary

- **MSC Nastran Performance**
  - ProLiant G7 servers provide 46% higher performance than best published results
  - Test illustrates Using more DMPs allow Nastran to scale
  - Time reduced as more nodes are being utilized for computation
  - Both Platform MPI and Intel MPI performs about the same
- **MSC Nastran Profiling**
  - Frequent used message sizes
    - Small message are used the most frequently
  - Frequent used MPI functions
    - MPI_Recv and MPI_Ssend
  - Heavy disk write access is seen throughout the test run
  - Tests shows that Nastran could benefit from better disk IO

Thank You

HPC Advisory Council