

# ANSYS Fluent 14 Performance Benchmark and Profiling

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

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

- <http://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
  - Enable better understanding of qualitative and quantitative physical phenomena in the flow 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**
  - Fluent performance benchmarking
  - Interconnect performance comparisons
  - MPI performance comparison
  - Understanding Fluent 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 1.5.3-3.10 InfiniBand SW stack
- **Mellanox ConnectX-3 VPI InfiniBand adapters**
- **Mellanox SwitchX SX6036 56Gb/s InfiniBand and 40G/s Ethernet Switch**
- **MPI (Vendor-provided): Intel MPI 4 U2, Open MPI 1.3.3, Platform MPI 8.1.2**
- **Application: ANSYS Fluent 14.0.0**
- **Benchmarks:**
  - Eddy\_417k – Reacting Flow with Eddy Dissipation Model (417K elements)

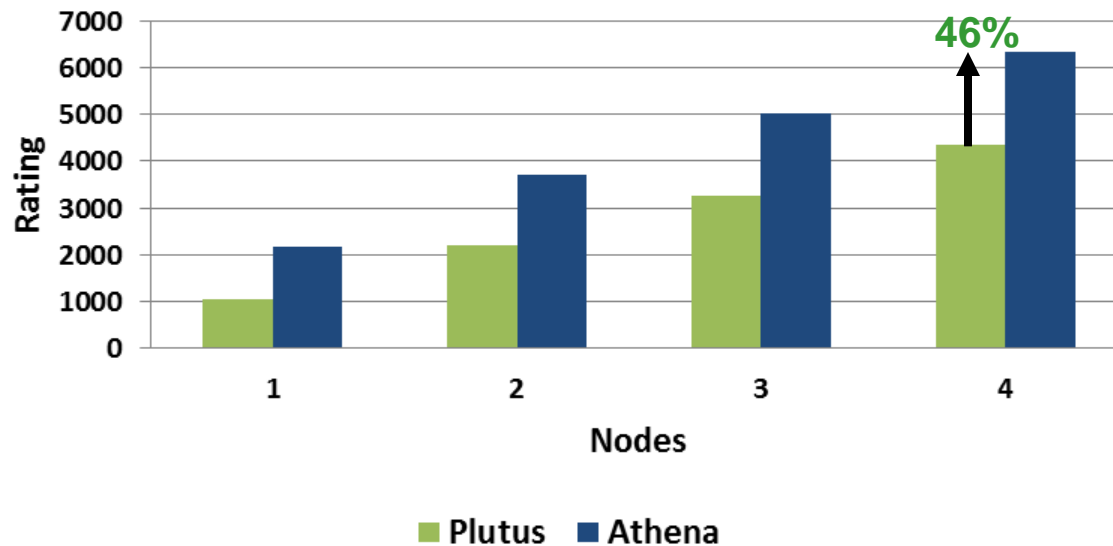
# About HP ProLiant SL230s Gen8

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

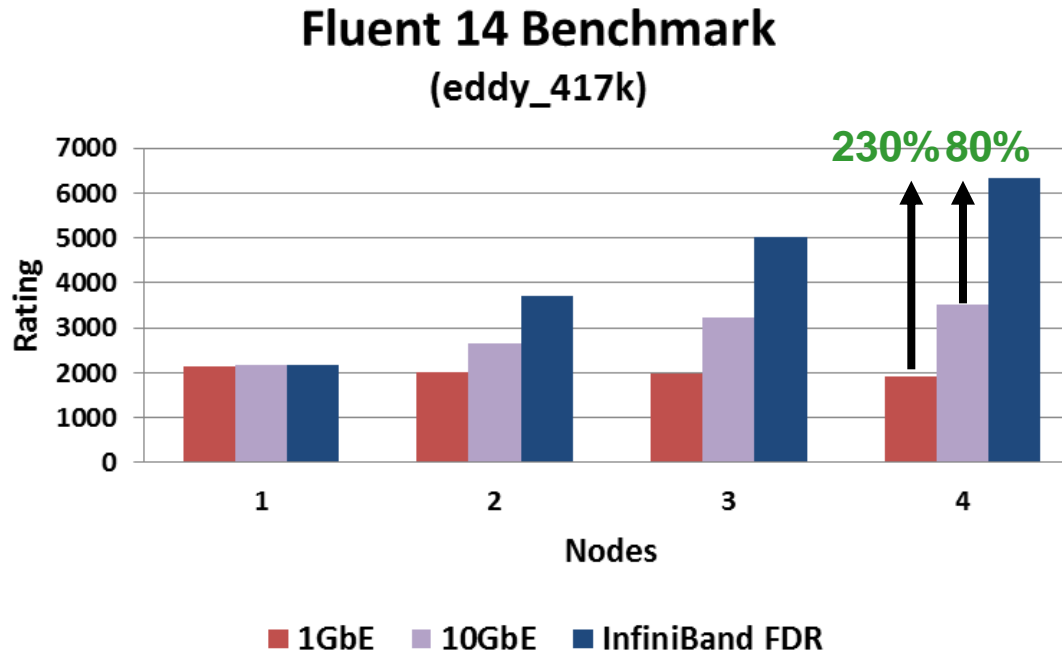


- **Intel E5-2680 (Sandy Bridge) cluster outperforms prior generations**
  - Performs 46% better than X5670 “Plutus” cluster nodes
- **System components used:**
  - Athena: 2-socket Intel E5-2680 @ 2.7GHz, 1600MHz DIMMs, FDR IB
  - Plutus: 2-socket Intel X5670 @ 2.93GHz, 1333MHz DIMMs, QDR IB

**Fluent 14 Benchmark  
(eddy\_417k)**

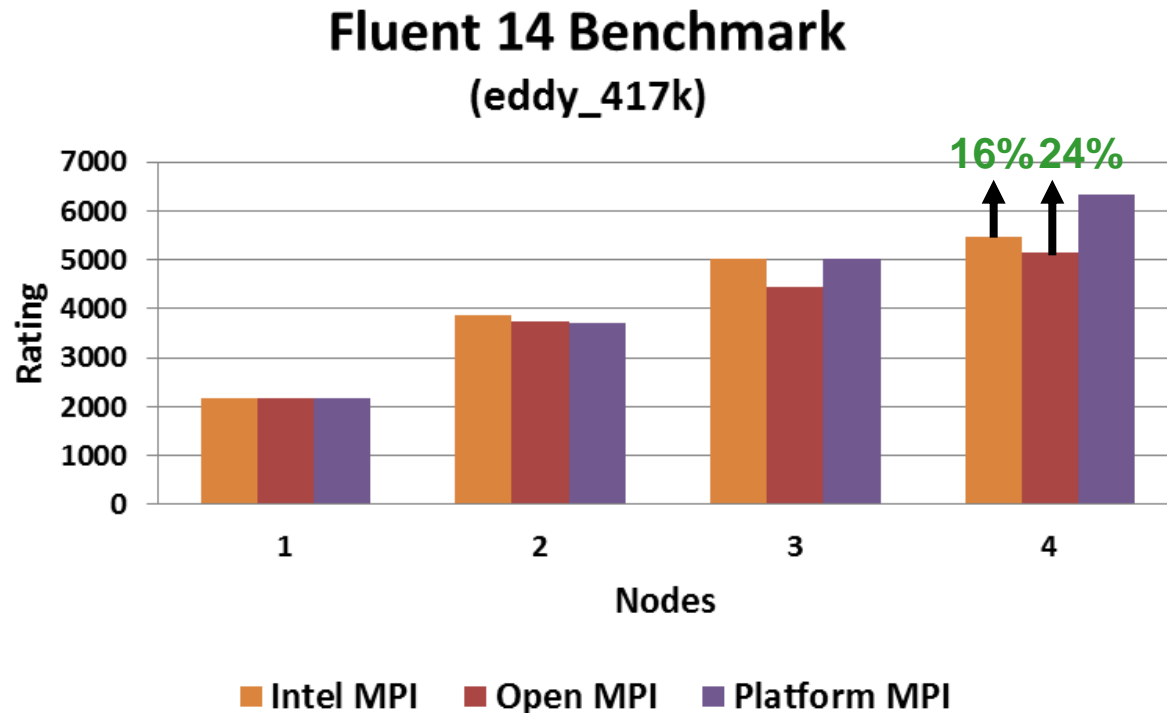


- **InfiniBand FDR is the most efficient inter-node communication for Fluent**
  - Outperforms 1GbE by 230% at 4 nodes
  - Outperforms 40GbE by 80% at 4 nodes





- **Platform MPI provides better scalability than Open MPI**
  - Fluent 14 provides Platform MPI as the default MPI option
  - Up to 24% better performance at 4 nodes
- **Default Fluent run script is used for all cases shown**
  - No other optimization flags were added




*Higher is better*

*InfiniBand FDR*


- **InfiniBand FDR reduces the time needed for communication**
  - InfiniBand FDR frees up more time for computation
  - Ethernet solutions consume from 74% to 86% of time in MPI communications

**Fluent 14 Profiling**  
(eddy\_417k, 4 nodes,  
1GbE)  
MPI/User Time Ratio




■ MPI time ■ User time

**Fluent 14 Profiling**  
(eddy\_417k, 4 nodes,  
10GbE)  
MPI/User Time Ratio



■ MPI time ■ User time

**Fluent 14 Profiling**  
(eddy\_417k, 4 nodes,  
InfiniBand FDR)  
MPI/User Time Ratio



■ MPI time ■ User time


*Higher is better*

*16 Processes/Node*

- **The most used MPI calls is MPI\_Iprobe**
  - Aside from MPI\_Iprobe, MPI\_Isend and Irecv are the next most used calls


*Showing All MPI calls*

**Fluent Profiling**  
(eddy\_417k)  
Number of MPI Calls



*Excluding MPI\_Iprobe*

**Fluent Profiling**  
(eddy\_417k)  
Number of MPI Calls

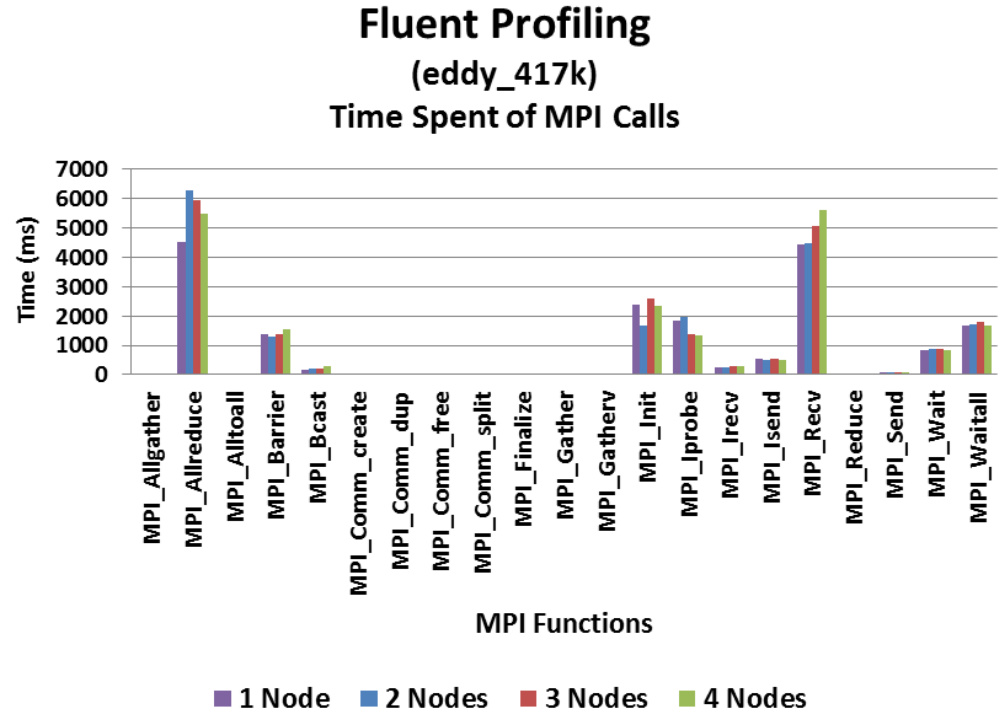
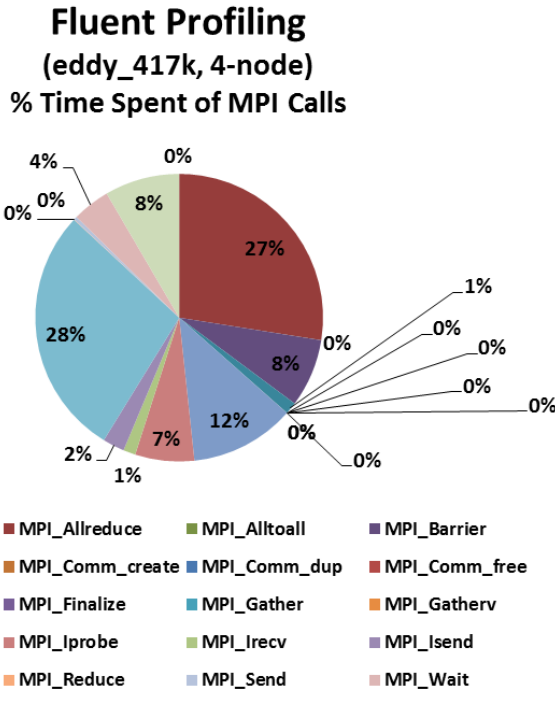


*Higher is better*

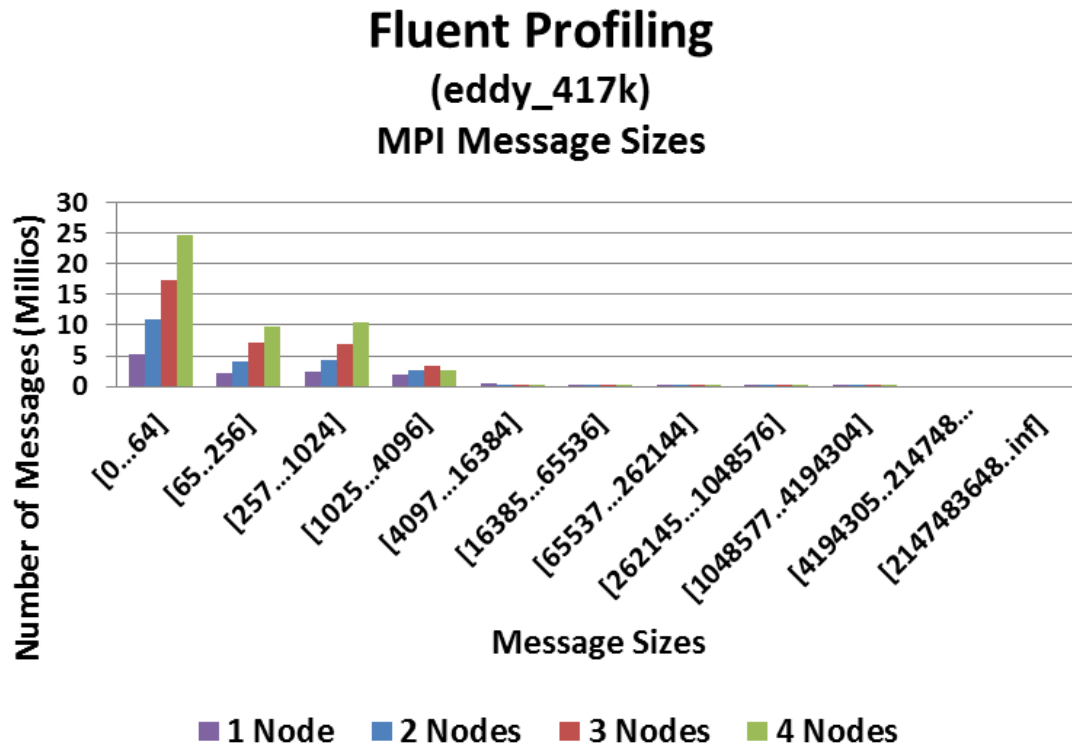
**16 Processes/Node**

# Fluent Profiling – MPI Communication Time

- Majority of MPI communication time is spent on MPI\_Recv
  - MPI\_Allreduce(27%), MPI\_Recv(28%)



- **Majority of messages are small messages**
  - Messages below 4KB are mostly used



- **HP ProLiant Gen8 servers delivers better performance than its predecessor**
  - ProLiant Gen8 equipped with Intel E5 series processes and InfiniBand FDR
  - Up to 46% higher performance than ProLiant G7 (running Intel Xeon X5670) when compared at 4 nodes
- **InfiniBand FDR is the most efficient inter-node communication for Fluent**
  - Outperforms 1GbE by 230% at 4 nodes
  - Outperforms 10GbE by 80% at 4 nodes
- **Fluent Profiling**
  - Platform MPI performs 24% better than Open MPI, and 16% better than Intel MPI
  - InfiniBand FDR reduces communication time; provides more time for computation
    - InfiniBand FDR consumes 53% of total time, versus 74-86% for Ethernet solutions
  - MPI:
    - Large MPI call volumes for testing non-blocking data transfers (MPI\_Iprobe, MPI\_Isend, MPI\_Irecv)
    - MPI time is spent mostly on MPI\_Recv and MPI\_Allreduce
    - Messages are concentrated in small messages, from 0B to 4KB

# 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