panasas

# Performance Study:
# FLUENT 12 and PanFS


***Stan Posey***
***Industry and Applications Market Development***
***Panasas, Fremont, CA, USA***

***Bill Loewe***
***Technical Staff Member, Applications Engineering***
***Panasas, Fremont, CA, USA***

# Panasas Company Overview

| | |
|---|---|
| Founded | 1999 By Prof. Garth Gibson, Co-Inventor of RAID |
| Technology | **Parallel File System and Parallel Storage Appliance** |
| Locations | US: HQ in Fremont, CA, USA<br>R&D centers in Pittsburgh & Minneapolis<br>EMEA: UK, DE, FR, IT, ES, BE, Russia<br>APAC: China, Japan, Korea, India, Australia |
| Customers | FCS October 2003, deployed at 200+ customers |
| Market Focus | Energy, Academia, Government, Life Sciences, Manufacturing, Finance |
| Alliances | ISVs: ANSYS, Microsoft Windows Compute Cluster Server<br>Resellers: DELL, sgi, Penguin Computing |
| Primary Investors | MDV Mohr Davidow Ventures, The Carlyle Group, Centennial Ventures, intel |

# Background on Parallel FLUENT Study

## Motivation

- Since 2006, Ansys and Panasas have jointly-invested in development of parallel I/O for release in FLUENT 12
- This study demonstrates benefits of Panasas parallel file system and parallel storage for FLUENT 12 vs. FLUENT 6.3 with tests for both capability and capacity computing
- Collaborators include Ansys and the University of Cambridge

## Considerations

- FLUENT is an application from ANSYS -- not a benchmark kernel
- The CFD models and tests and relevant to customer practice
- This was run on a production system at customer U of Cambridge
- The results were validated by U of Cambridge and ANSYS

# Description of Darwin at U of Cambridge

## University of Cambridge

UNIVERSITY OF CAMBRIDGE

**HPC Service, Darwin Supercomputer**

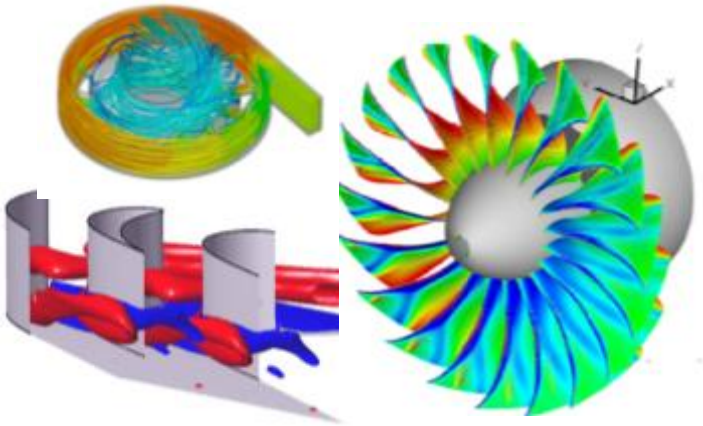

### Darwin Supercomputer Computational Units

- Nine repeating units, each consists of 64 nodes (2 racks) providing 256 cores each, 2340 cores total
- All nodes within a CU connected to a full bisectional bandwidth Infiniband 900 MB/s, MPI latency of 2 μs

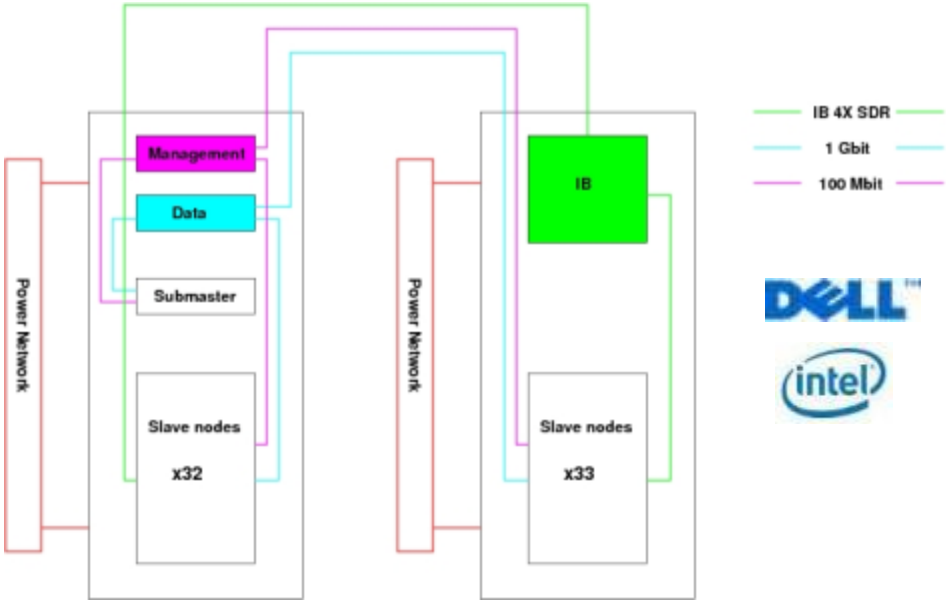


**Source:** http://www.hpc.cam.ac.uk

# Details of the FLUENT 111M Cell Model

## Unsteady external aero for 111 MM cell truck; 5 time steps with 100 iterations, and a single *.dat* file write

| | |
|---|---|
| **Number of cells** | **111,091,452** |
| **Solver** | **PBNS, DES, Unsteady** |
| **Iterations** | **5 time steps, 100 total iters - data save after last iteration** |

**Output size:**

| | | |
|---|---|---|
| **FLUENT v6.3** | **(serial I/O; size of .dat file)** | **14,808 MB** |
| **FLUENT v12** | **(serial I/O; size of .dat file)** | **16,145 MB** |
| **FLUENT v12** | **(parallel I/O; size of .pdat file)** | **19, 683 MB** |

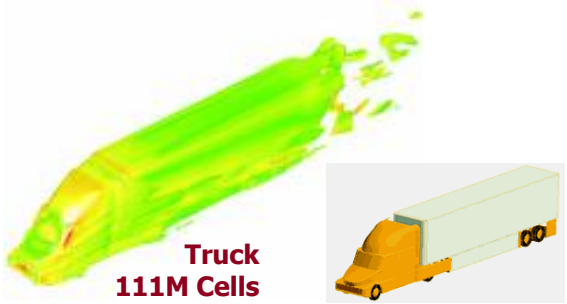


**DARWIN 585 nodes; 2340 cores**

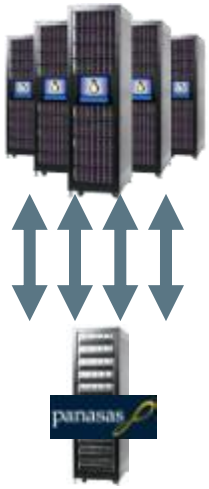


**Panasas: 4 Shelves, 20 TB**

| **Univ of Cambridge DARWIN Cluster** |
|---|
| **Location:** University of Cambridge http://www.hpc.cam.ac.uk |
| **Vendor:** Dell ; 585 nodes; 2340 cores; 8 GB per node; 4.6 TB total mem |
| **CPU:** Intel Xeon (Woodcrest ) DC, 3.0 GHz / 4MB L2 cache |
| **Interconnect:** InfiniPath QLE7140 SDR HCAs; Silverstorm 9080 and 9240 switches, |
| **File System:** Panasas PanFS, 4 shelves, 20 TB capacity |
| **Operating System:** Scientific Linux CERN SLC release 4.6 |

# This Study is a Partial CFD Simulation

## Unsteady CFD Simulation Schematic and Typical I/O Profile

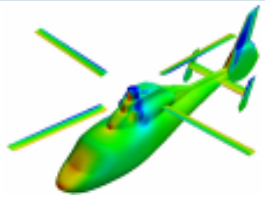

**The Focus of the FLUENT study is only a sub-set of a full unsteady CFD simulation:**

- Read once
- Compute 5 time steps (100 iters)
- Write once **(but full simulation has multiple writes)**

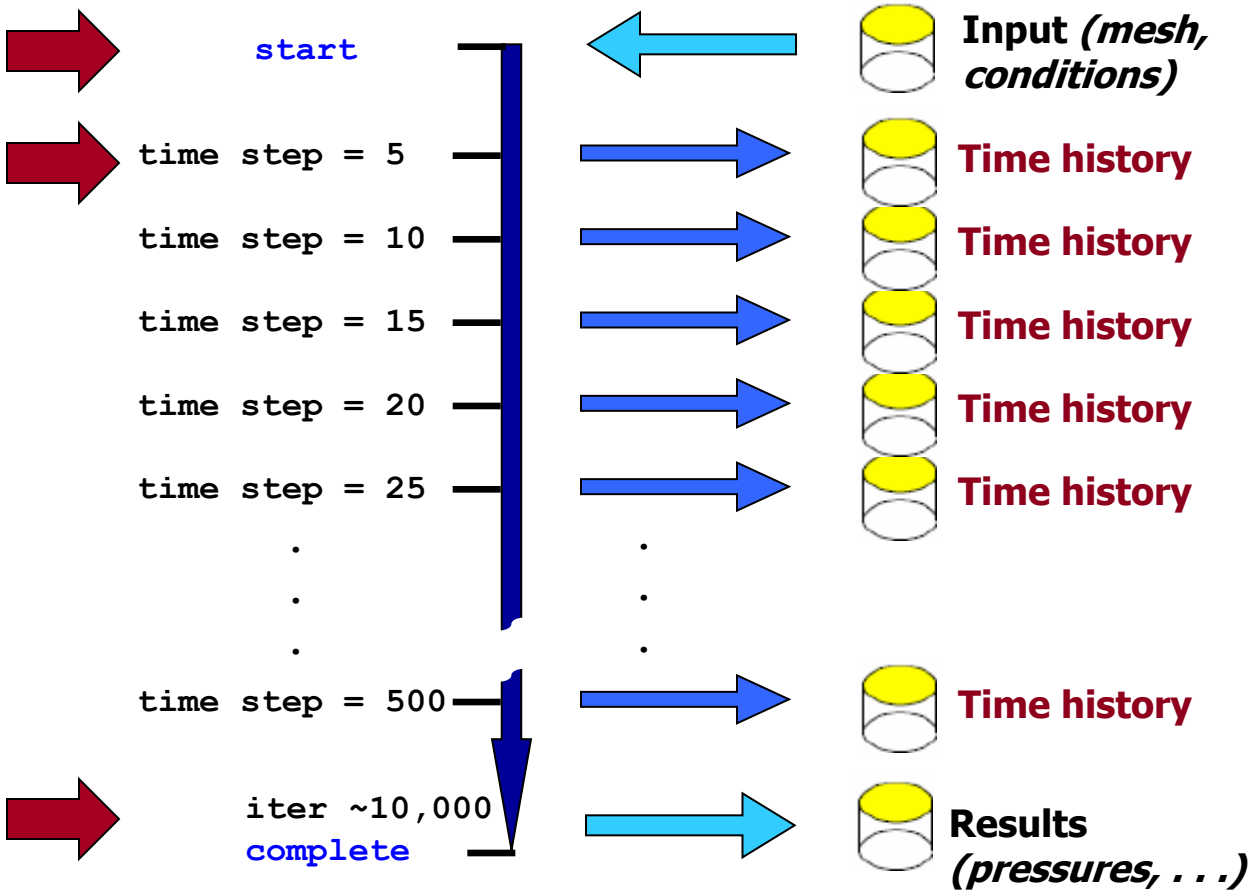

# FLUENT 12 and New Parallel I/O Scheme

panasas

## Panasas and ANSYS Alliance Has Produced Parallel I/O for FLUENT 12

ANSYS

**Serial I/O Scheme** **Improvement** **Parallel I/O Scheme**

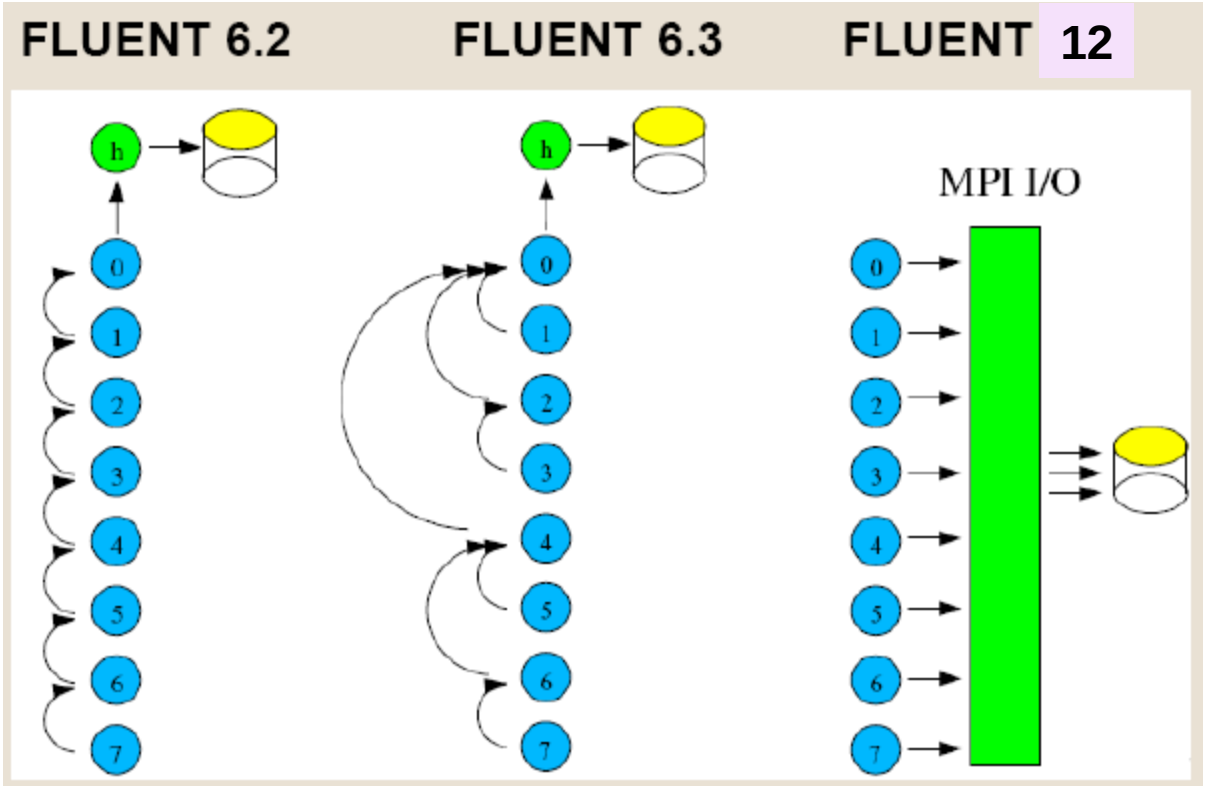


**FLUENT 12:** Offers support for PanFS and most other parallel file systems

Beta available today

Source: Barb Hutchings Presentation at SC07, Nov 2007, Reno, NV

# Parallel FLUENT 12 and 2x Improvement

panasas

## FLUENT Comparison of PanFS vs. NFS on University of Cambridge Cluster

UNIVERSITY OF CAMBRIDGE

ANSYS

panasas

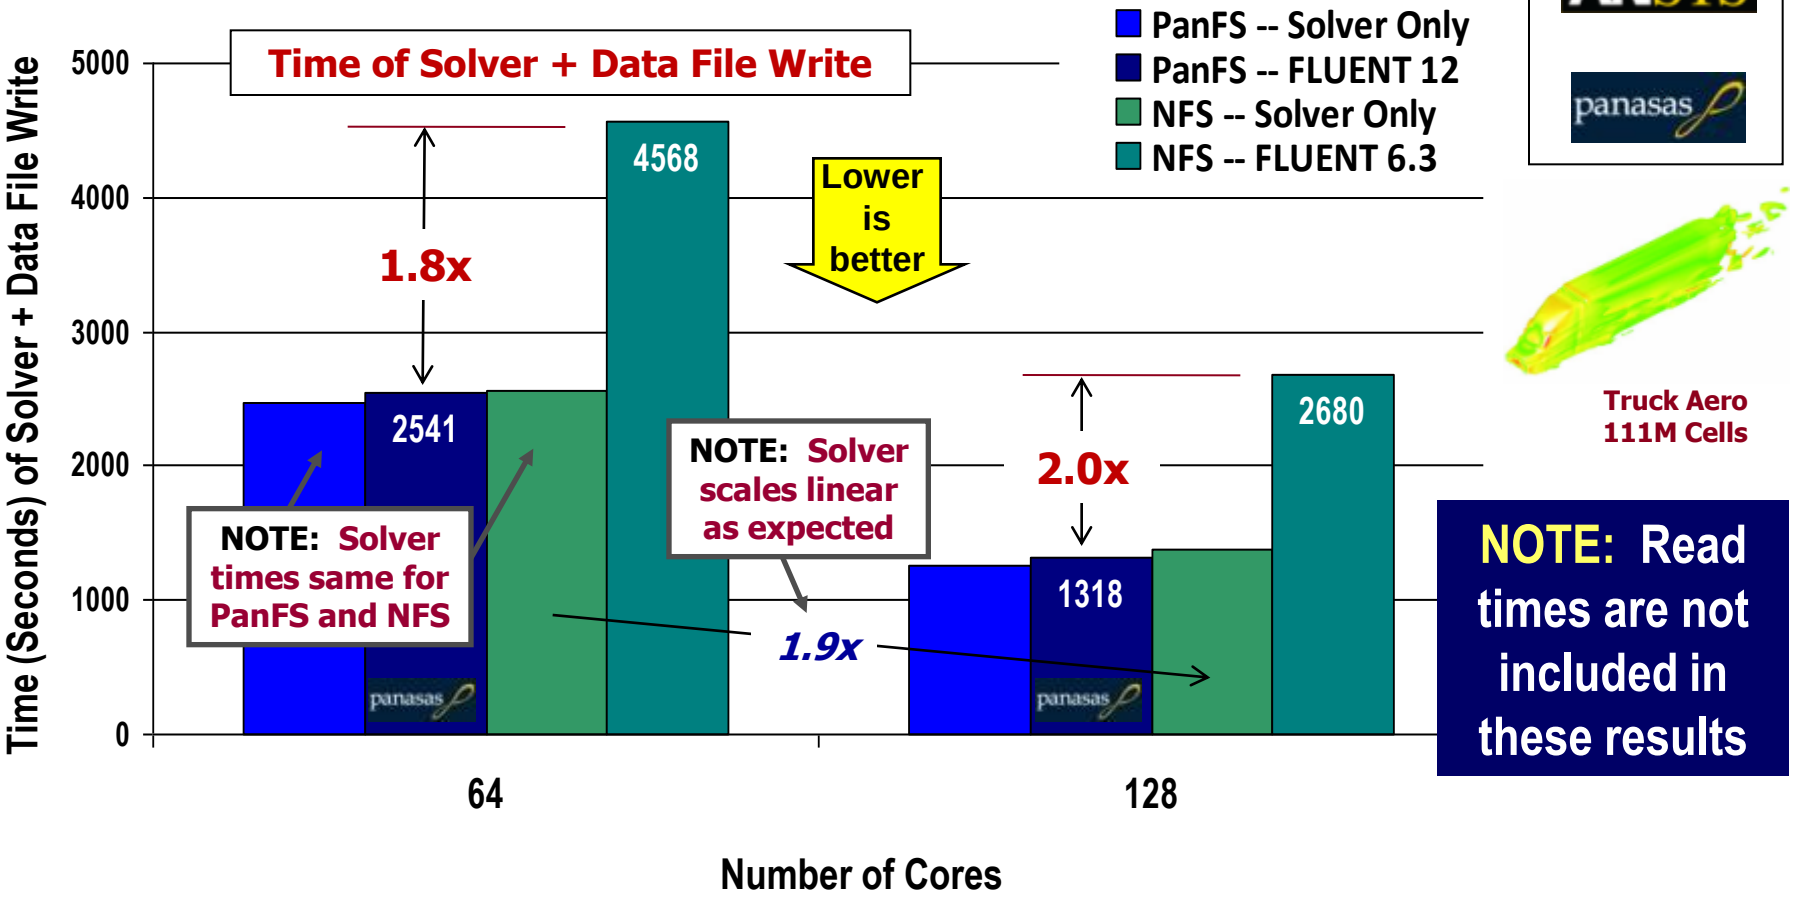




Truck Aero
111M Cells

**NOTE: Read times are not included in these results**

# Scalability of Solver + Data File Write

**FLUENT Comparison of PanFS vs. NFS on University of Cambridge Cluster**

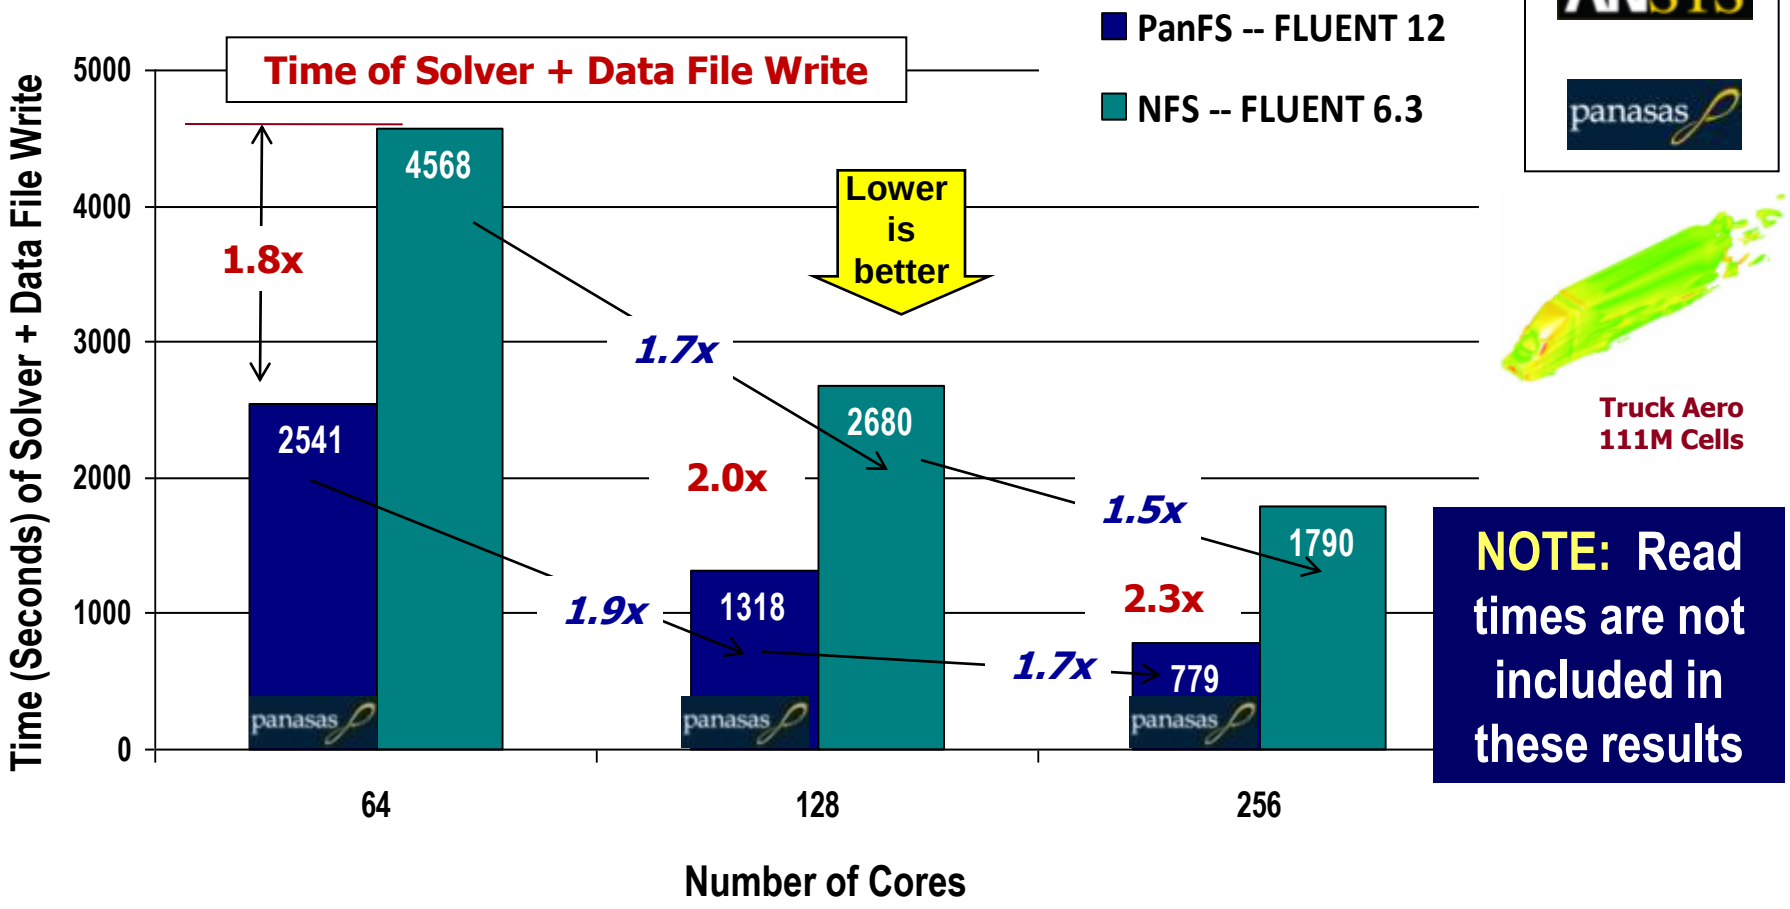




Truck Aero
111M Cells

**NOTE: Read times are not included in these results**

# Performance of Data File Write in MB/s

**FLUENT Comparison of PanFS vs. NFS on University of Cambridge Cluster**

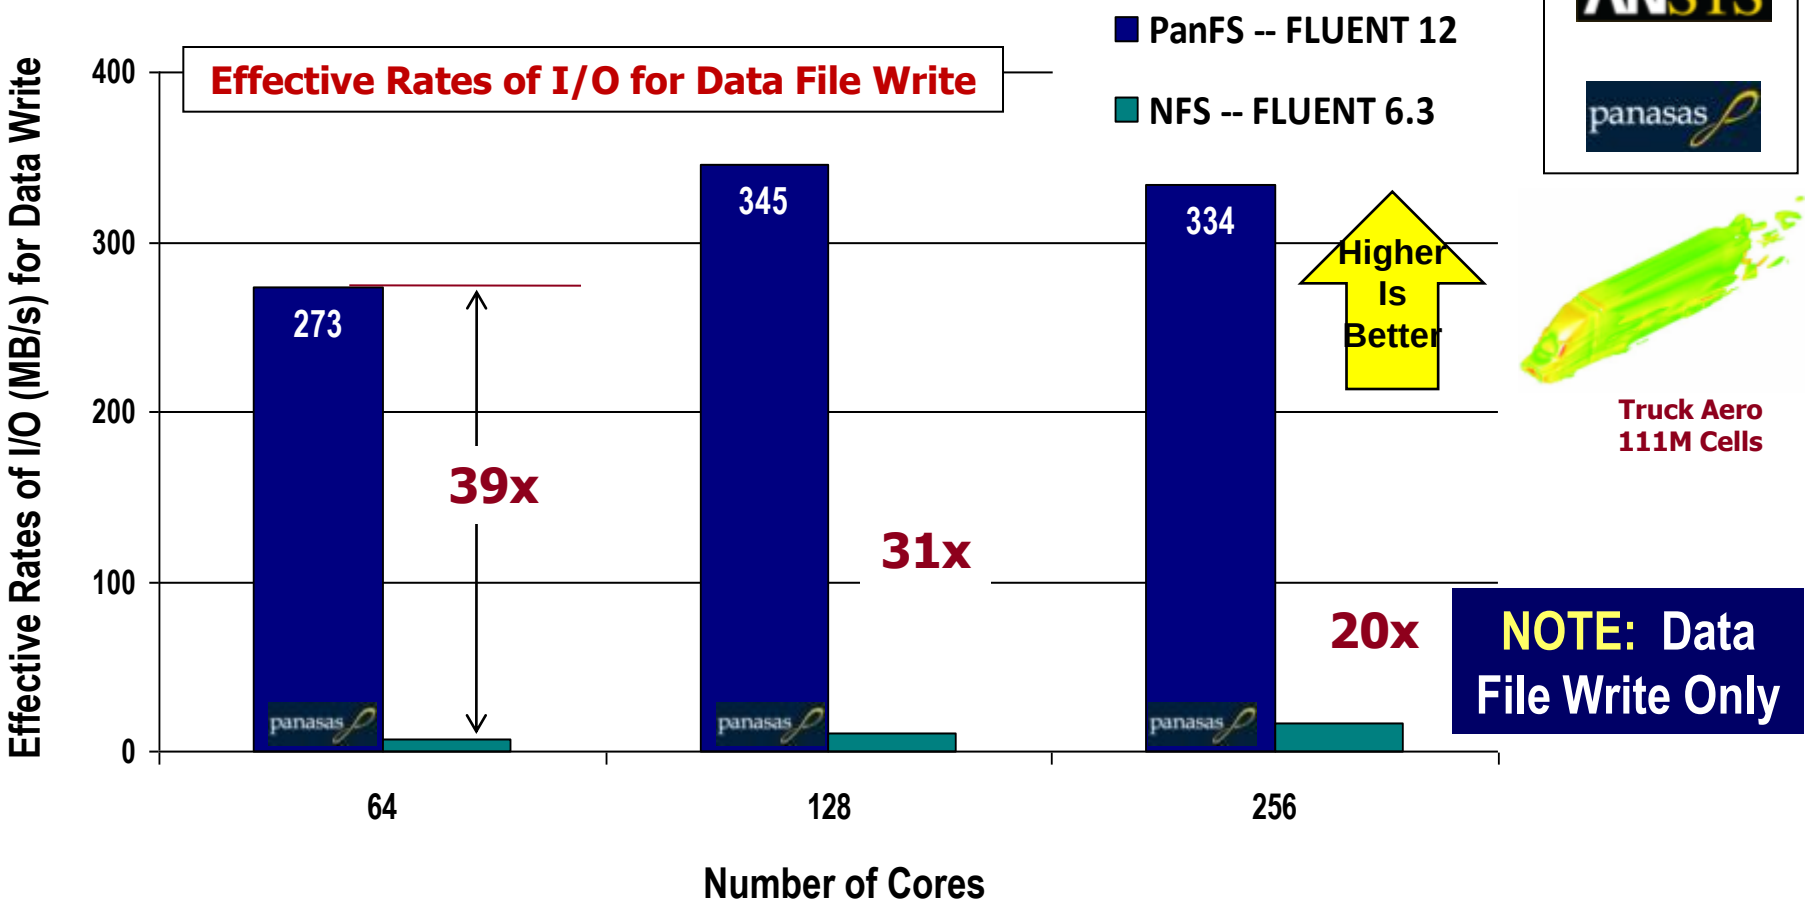

# Observations From FLUENT 12 Study

panasas

- FLUENT 6.3 End-User Challenges
  - Large production cases may not scale effectively and efficiently on a large cluster (>64 cores) for read and write operations owing to serial I/O
  - The use of frequent checkpoints for very large steady-state cases, and/or large unsteady simulations (multiple writes) is impractical with serial I/O
- FLUENT 12 and Panasas Solution
  - The Panasas parallel file system and storage, combined with parallel I/O of FLUENT 12 scales I/O and therefore the overall FLUENT simulation
  - Use of PanFS for the 111M cell case at 64-way provides a nearly 2x increase in FLUENT utilization for the same software license $'s spent
  - Such capability enables FLUENT users to develop more advanced CFD models (more transient vs. steady, LES, etc.) with confidence in scalability

# Two Measures of FLUENT Performance

- A single wide-parallel job vs. multiples of less-parallel jobs
  - Often referenced in HPC industry as capability vs. capacity computing
  - Both are important, but capacity computing more common in practice
  - Example: design optimization based on capacity, impractical with capability

**Panasas scales I/O for the large single CFD job, and provides parallel data access (vs. serial NFS) for multi-job scenarios**

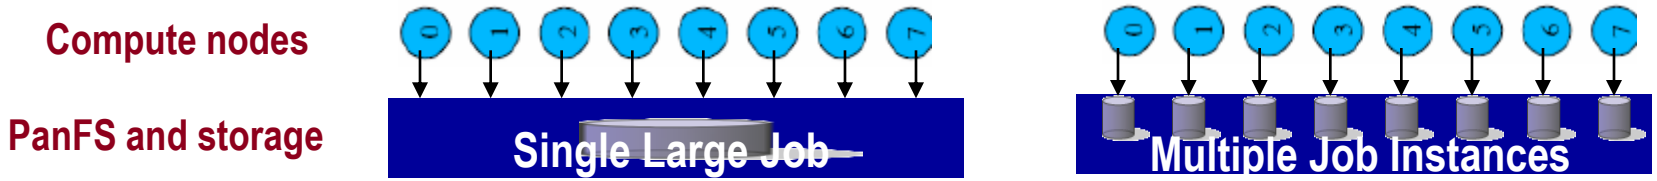


- A multi-job test was developed with the Truck model at 14M cells:
  - The same Truck model with a coarsened mesh from 111M to 14M cells
  - Launched 8 times (8 copies) each using 16 cores for a total 128 cores
  - PanFS is parallel, NFS has single data path for during 8 solution writes

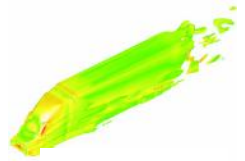
Truck Aero
14M Cells

# Details of the FLUENT 14M Cell Model

## Unsteady external aero for 14 MM cell truck; 5 time steps with 100 iterations, and a single *.dat* file write

| | |
|---|---|
| **Number of cells** | **13,839,118** |
| **Solver** | **PBNS, DES, Unsteady** |
| **Iterations** | **5 time steps, 100 total iters - data save after last iteration** |

**Output size:**

| | | |
|---|---|---|
| **FLUENT v6.3** | **(serial I/O; size of .dat file)** | **2,466 MB** |
| **FLUENT v12** | **(serial I/O; size of .dat file)** | **2,467 MB** |
| **FLUENT v12** | **(parallel I/O; size of .pdat file)** | **2,898 MB** |


**DARWIN 585 nodes; 2340 cores**

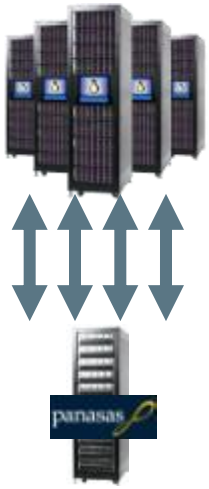

**Panasas: 4 Shelves, 20 TB**

### Univ of Cambridge DARWIN Cluster

| |
|---|
| **Location:** University of Cambridge http://www.hpc.cam.ac.uk |
| **Vendor:** Dell ; 585 nodes; 2340 cores; 8 GB per node; 4.6 TB total mem |
| **CPU:** Intel Xeon (Woodcrest )DC, 3.0 GHz / 4MB L2 cache |
| **Interconnect:** InfiniPath QLE7140 SDR HCAs; Silverstorm 9080 and 9240 switches, |
| **File System:** Panasas PanFS, 4 shelves, 20 TB capacity |
| **Operating System:** Scientific Linux CERN SLC release 4.6 |

# Panasas Benefit with Multi-Job FLUENT

## FLUENT Comparison of PanFS vs. NFS on University of Cambridge Cluster

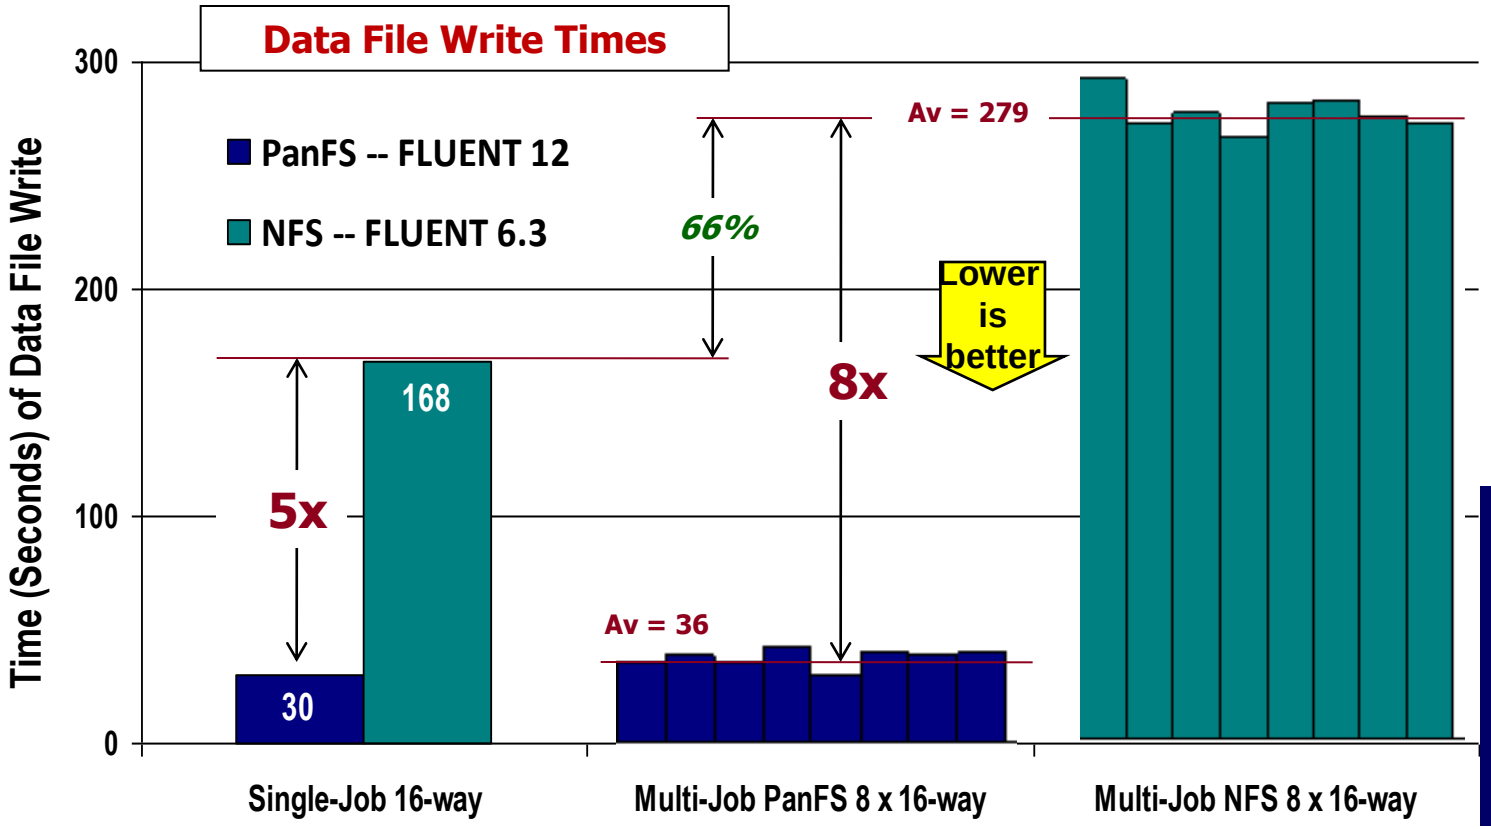


UNIVERSITY OF CAMBRIDGE

ANSYS

panasas

Truck Aero
14M Cells

**NOTE:** NFS degrades 66% on average for the multi-job scenario test

# Performance of Data File Write in MB/s

panasas

## FLUENT Comparison of PanFS vs. NFS on University of Cambridge Cluster

UNIVERSITY OF CAMBRIDGE

ANSYS

panasas

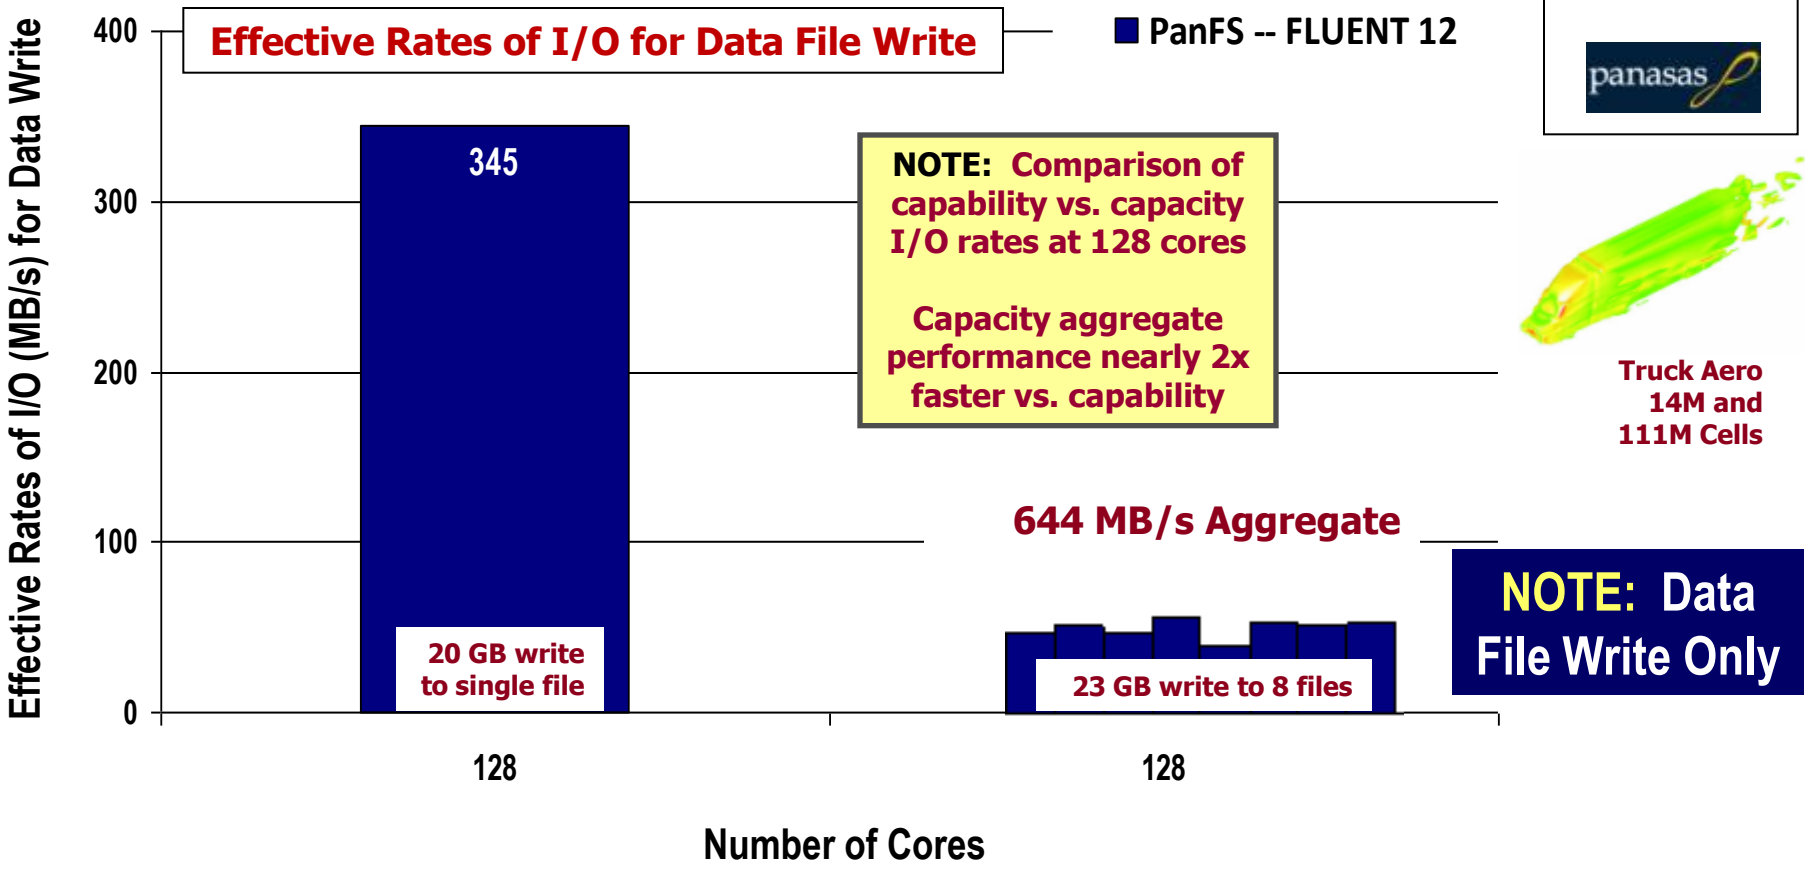


**NOTE: Comparison of capability vs. capacity I/O rates at 128 cores**

**Capacity aggregate performance nearly 2x faster vs. capability**



Truck Aero 14M and 111M Cells

**NOTE: Data File Write Only**

# Contributors to the Study

**University of Cambridge**

- Dr. Paul Calleja, Director, HPCS
- Dr. Stuart Rankin, Lead System Manager, HPCS

**ANSYS**

- Dr. Prasad Alavilli, FLUENT and CFX Development
- Ms. Barbara Hutchings, Director of Technology Alliances

**Panasas**

- Mr. Derek Burke, Director of Marketing, Panasas EMEA
- Ms. Michelle Cheng, Director of Global Alliances

# Thank You for This Opportunity

panasas

**RESOURCES:**

- **Questions can be directed to the Panasas email addresses below**
- **The 111M cell truck model is public and available from Ansys**

  **http://www.fluent.com/software/fluent/fl6bench/fl6bench_6.3/problems/truck_111m.htm**
- **FLUENT log files of all jobs are available upon request to Panasas**

**Stan Posey**
sposey@panasas.com

**Bill Loewe**
bloewe@panasas.com