



Performance Study: **STAR-CD v4 on PanFS**



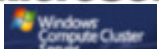







*Stan Posey
Industry and Applications Market Development
Panasas, Fremont, CA, USA*

*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 <u>Appliance</u>						
Locations	<p>US: HQ in Fremont, CA, USA R&D centers in Pittsburgh & Minneapolis</p> <p>EMEA: UK, DE, FR, IT, ES, BE, Russia</p> <p>APAC: China, Japan, Korea, India, Australia</p>						
Customers	FCS October 2003, deployed at 200+ customers						
Market Focus	<table border="0"> <tr> <td>Energy</td> <td>Academia</td> </tr> <tr> <td>Government</td> <td>Life Sciences</td> </tr> <tr> <td>Manufacturing</td> <td>Finance</td> </tr> </table>	Energy	Academia	Government	Life Sciences	Manufacturing	Finance
Energy	Academia						
Government	Life Sciences						
Manufacturing	Finance						
Alliances	<p>ISVs:   </p> <p>Resellers:   </p>						
Primary Investors	   						

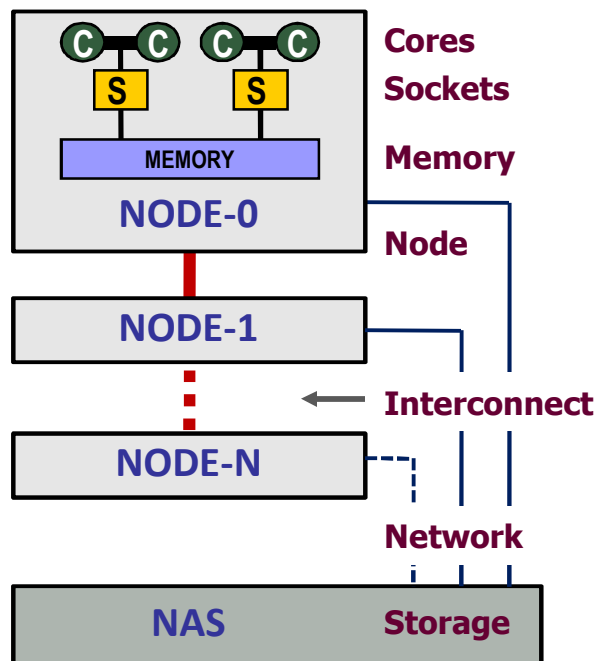
HPC Characterization of STAR-CD (I)



Like most all parallel CFD, a STAR-CD job contains a mix of compute tasks that each require specific performance attributes of an HPC system:

- Numerical Operations: typically equations solvers and other modeling calculations
- Communication Operations: partition boundary information “passed” between cores
- Read and Write Operations: data file i/o before/during/after computations

Schematic of HPC System Stack



STAR-CD Compute Task

Numerical Operations

Communications (MPI)

Read/Write Operations

Performance Attribute

Fast CPU architectures to speed-up equation solvers operating on each partition

Low-latency interconnects and MPI system software to minimize communications overhead between partitions for higher levels of scalability

Parallel file system with NAS to ensure concurrent reads and writes that scale the I/O

What Does This Characterization Mean for File Systems and Storage?

- File systems and storage affect performance of read/write operations only
- All other operations (numerical and communications , for example in STAR equation solvers) are NOT affected by the choice of file system and storage
- Therefore, computational profiles of STAR jobs that spend a large % of their total time in read/write operations will benefit from a parallel file system
- Examples of STAR-CD jobs with a large % of write operations:
 - Large (> 30M cells) parallel steady models with large output data files
 - Any meaningful size, parallel transient model (e.g. URANS, DES, LES)
 - Any moving mesh model, multi-phase VOF, frequent checkpoints . . .
 - A mix of multiple STAR jobs with concurrent I/O requests to a file system

NOTE: While numerical and communication operations for a STAR job “bind” to specific sets of nodes and interconnects in order to minimize job conflict for these resources, the file system and storage are a shared resource that must manage all concurrent (and competing) requests for read/write operations to the file system

Motivation

- Since 2007, CD-adapco and Panasas have jointly-invested in the development of parallel I/O for STAR-CD v4
- This study demonstrates benefits of Panasas parallel file system and parallel storage for STAR-CD v4
- Collaborators include CD-adapco and Intel Corporation



Considerations

- STAR-CD is an application from CD-adapco -- not a benchmark kernel
- The CFD model is large and relevant to customer practice
- Panasas storage is certified for Intel Cluster Ready (ICR) www.panasas.com
- This was run on an ICR system at the Intel HPC benchmark center
- The results were reviewed and validated by Intel and CD-adapco

Intel HPC Data Center Based on Panasas

Source: HPC Software @ Intel by Dr. Paresh Pattani, SC07, 12 Nov 07, Reno NV

Panasas and Intel HPC:
Unique relationship gives
Panasas certification on
wide range of applications


- CAE (Linux*)**
 - ABAQUS Standard, Explicit
 - Ansys (ANSYS)
 - CFD++ (Metacomp)
 - CFD-ACE+ (ESI)
 - Feko* (EMSS)
 - Fluent* (Fluent)
 - FIRE* (AVL)
 - HyperMesh* (Altair)
 - LS-Dyna* (LSTC)
 - Mafia* (CST)
 - MSC.Nastran* (MSC Software)
 - NX Nastran* (UGS)
 - PAM-Crash* (ESI)
 - PowerFLOW* (EXA)
 - Star-CD* (CD/Adapco)

Key Applications Enabled on 64-bit Intel® Xeon® Processor

- CAE (Linux*)**
 - ABAQUS Standard, Explicit
 - Ansys (ANSYS)
 - CFD++ (Metacomp)
 - CFD-ACE+ (ESI)
 - Feko* (EMSS)
 - Fluent* (Fluent)
 - FIRE* (AVL)
 - HyperMesh* (Altair)
 - LS-Dyna* (LSTC)
 - Mafia* (CST)
 - MSC.Nastran* (MSC Software)
 - NX Nastran* (UGS)
 - PAM-Crash* (ESI)
 - PowerFLOW* (EXA)
 - Star-CD* (CD/Adapco)
- Energy (Linux)**
 - Eclipse* (Schlumberger)
 - Petrel / GigaM2 (Schlumberger)
 - Omega2 (Western Geac)
 - VIP / Nexus (Landmark Graphics)
 - Geoprobe (Landmark Graphics)
 - Earthx / SeisSpace (Landmark Graphics)
 - Geodapt* (Paradigm Geo)
 - Focus* (Paradigm Geo)
 - GeoCluster (CBG)
- Life Sciences (Linux*)**
 - BLAST
 - Gaussian
 - AMBER
 - GAMESS
 - Gromacs
 - HMMER
 - NAMD
- NWS (Linux*)**
 - MYS
 - WRF
 - COSMB
 - POP
 - CAM
 - Aladin
 - HirLAM
 - UM
- Financial Services (Linux*)**
 - BisWatch (Algorithmic)
 - RMDS (Reuters)
 - Thomson Financial
 - Credit, Intelligents (Sungard)
 - OPS (Morgan-Stanley)
- DCC (Linux / Windows*)**
 - Alias - Maya
 - Discreet - 3ds max
 - mental images - mental ray
 - Pixar RenderMan
- CAD (Microsoft Windows*)**
 - Pro/ENGINEER* (PTC)
 - NX* (UGS)
 - Parasolid* (UGS)
 - I2 Vis* (UGS)

Workstation applications

Comprehensive Set of Applications



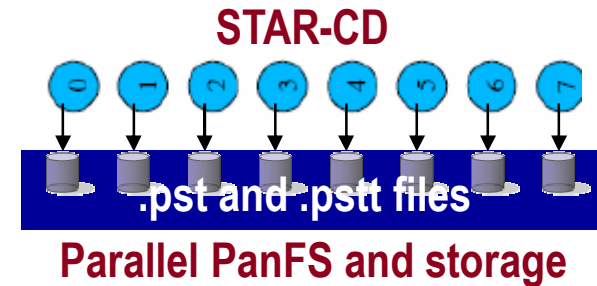
* Other brands and names may be claimed as the property of others. Intel Confidential

Joint Investments in Parallel I/O for STAR-CD and STAR-CCM+



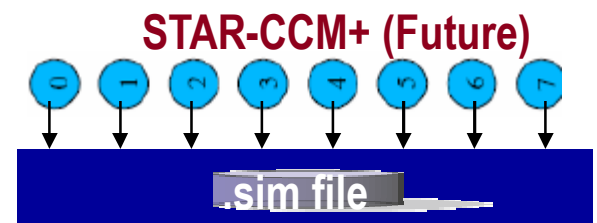
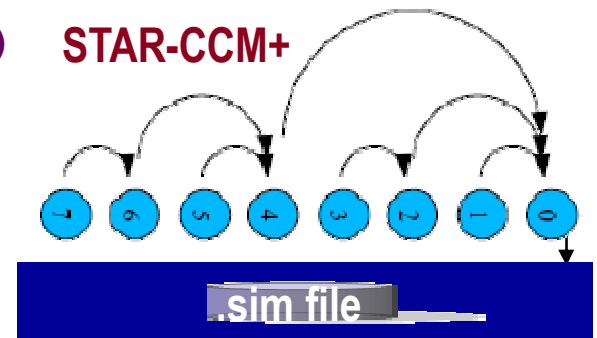
STAR-CD

- Parallel writes in v3.26, serial file merge at job completion
- Parallel writes in v4.06 without any file merge operations
- Applications with most benefit: Large (> 30M) cells steady; any URANS; LES; moving mesh (combustion); VOF (free surface); multi-phase; and weakly-coupled FSI (e.g. Abaqus)



STAR-CCM+

- Good performance with efficient serial I/O scheme today
- Stated plans for parallel I/O, 2009 roadmap under review
- Applications with most benefit: Large (> 30M) cells steady; any large-scale aerodynamics, aeroacoustics; CFD model parameterization with multiple-jobs making I/O requests to a shared file system



Panasas and Intel STAR-CD v4 Study




Steady external aero for 20 MM cell vehicle; 500 iterations (non-converged) and solution writes at every 10 iterations



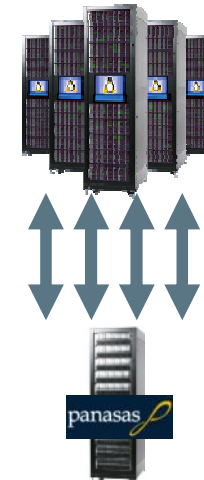
Number of cells	19,921,786
Solver	CGS, Steady
Iterations	500 total iterations - data save after every 10 iters
Each solution output (50 total)	~1,500 MB

**A-Class
20M Cells**



Intel "ENDEAVOR" Xeon® 
Nodes: 256 x 2 Sockets x 4 Cores = Total Cores: 2048
Location: DuPont, WA
CPU: Harpertown Xeon QC 2.8 GHz / 12MB L2 cache
FSB: 1600 MHz, IB Interconnect DDR
File Systems -- Panasas,: 7 shelves, 35 TB storage; NFS: Dell 2850 File Server, 6 x 146 GB SCSI drives, RAID 5
FS Connectivity: Gig Ethernet, 4 bonded links per shelf, 2.8 GB / sec peak, 2.5 GB./ sec measured

ENDEAVOR, 2048 cores



Panasas: 7 Shelves, 35 TB

This Study is a Partial CFD Simulation

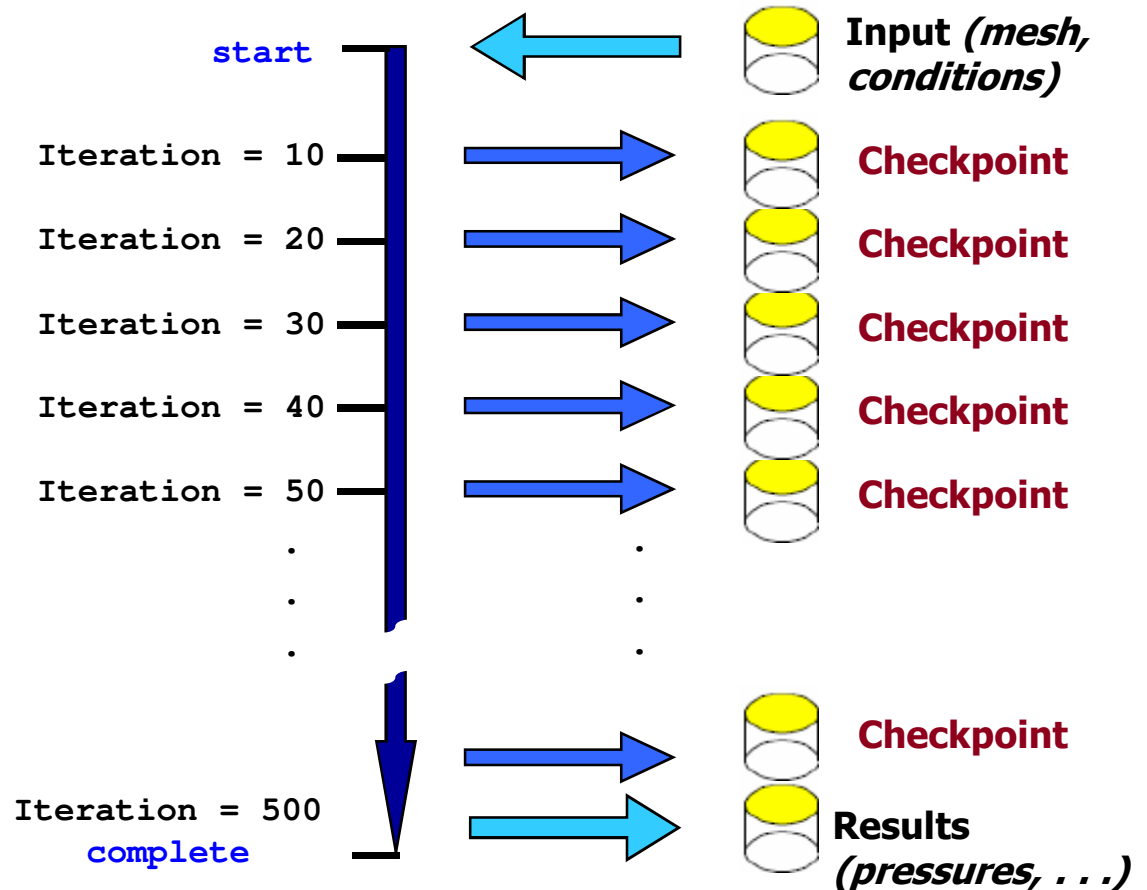


CFD Simulation Schematic and Typical I/O Profile



The focus of this STAR-CD study is only a sub-set of a full steady state CFD simulation:

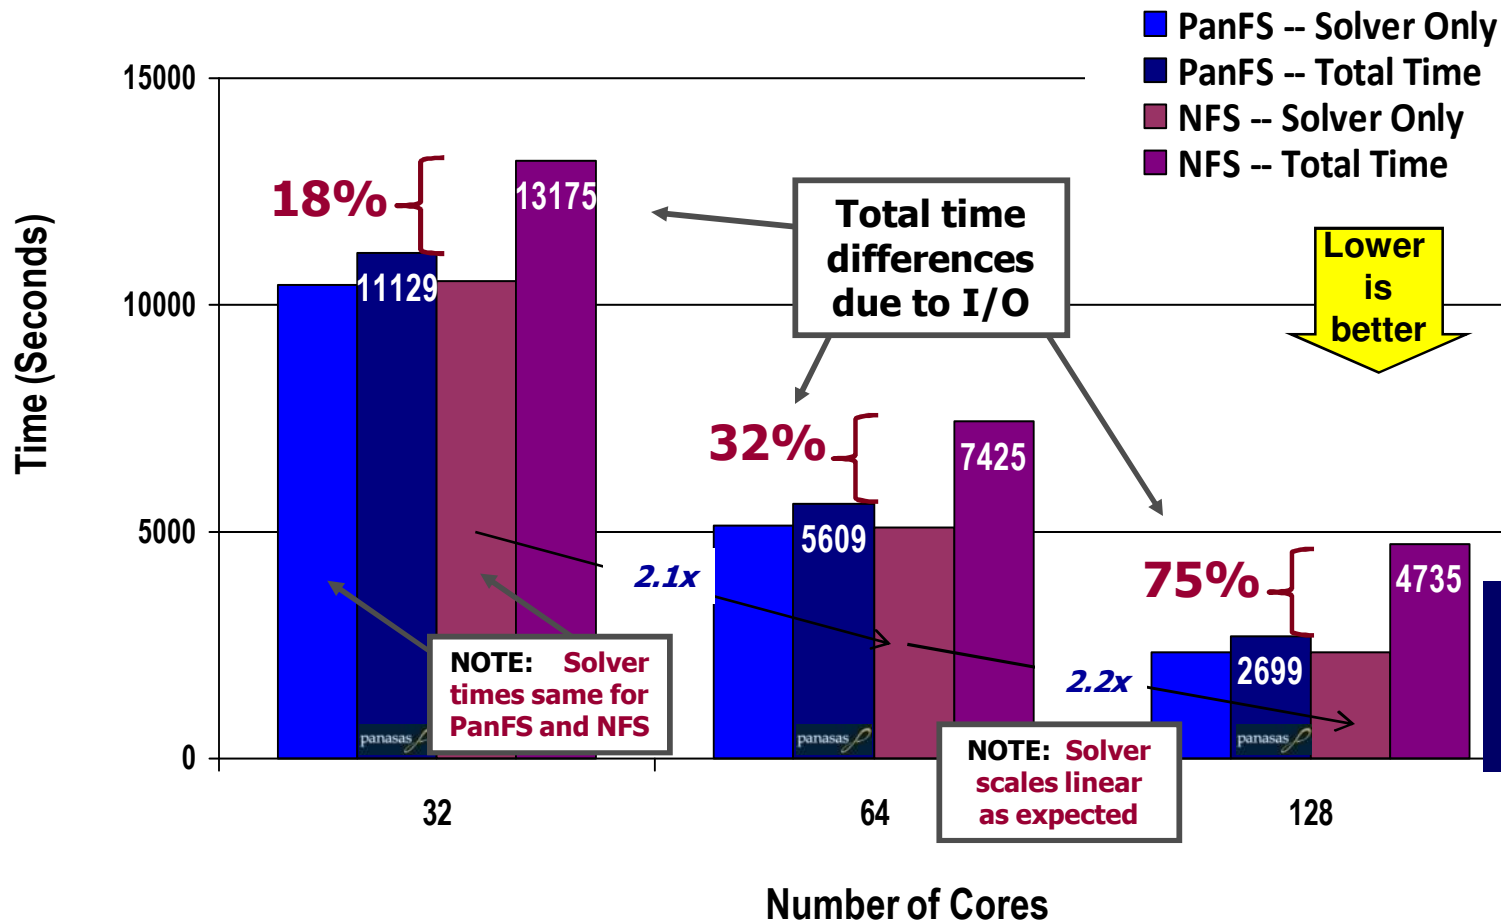
- Read once
- Compute 10 iters
- Write
- Compute 10 iters
- Write
- ⋮
- Stop at 500 iters (total of 50 writes)



STAR-CD v4 Total Time Comparisons



STAR-CD v4.06: Comparison of PanFS vs. NFS on Intel Cluster



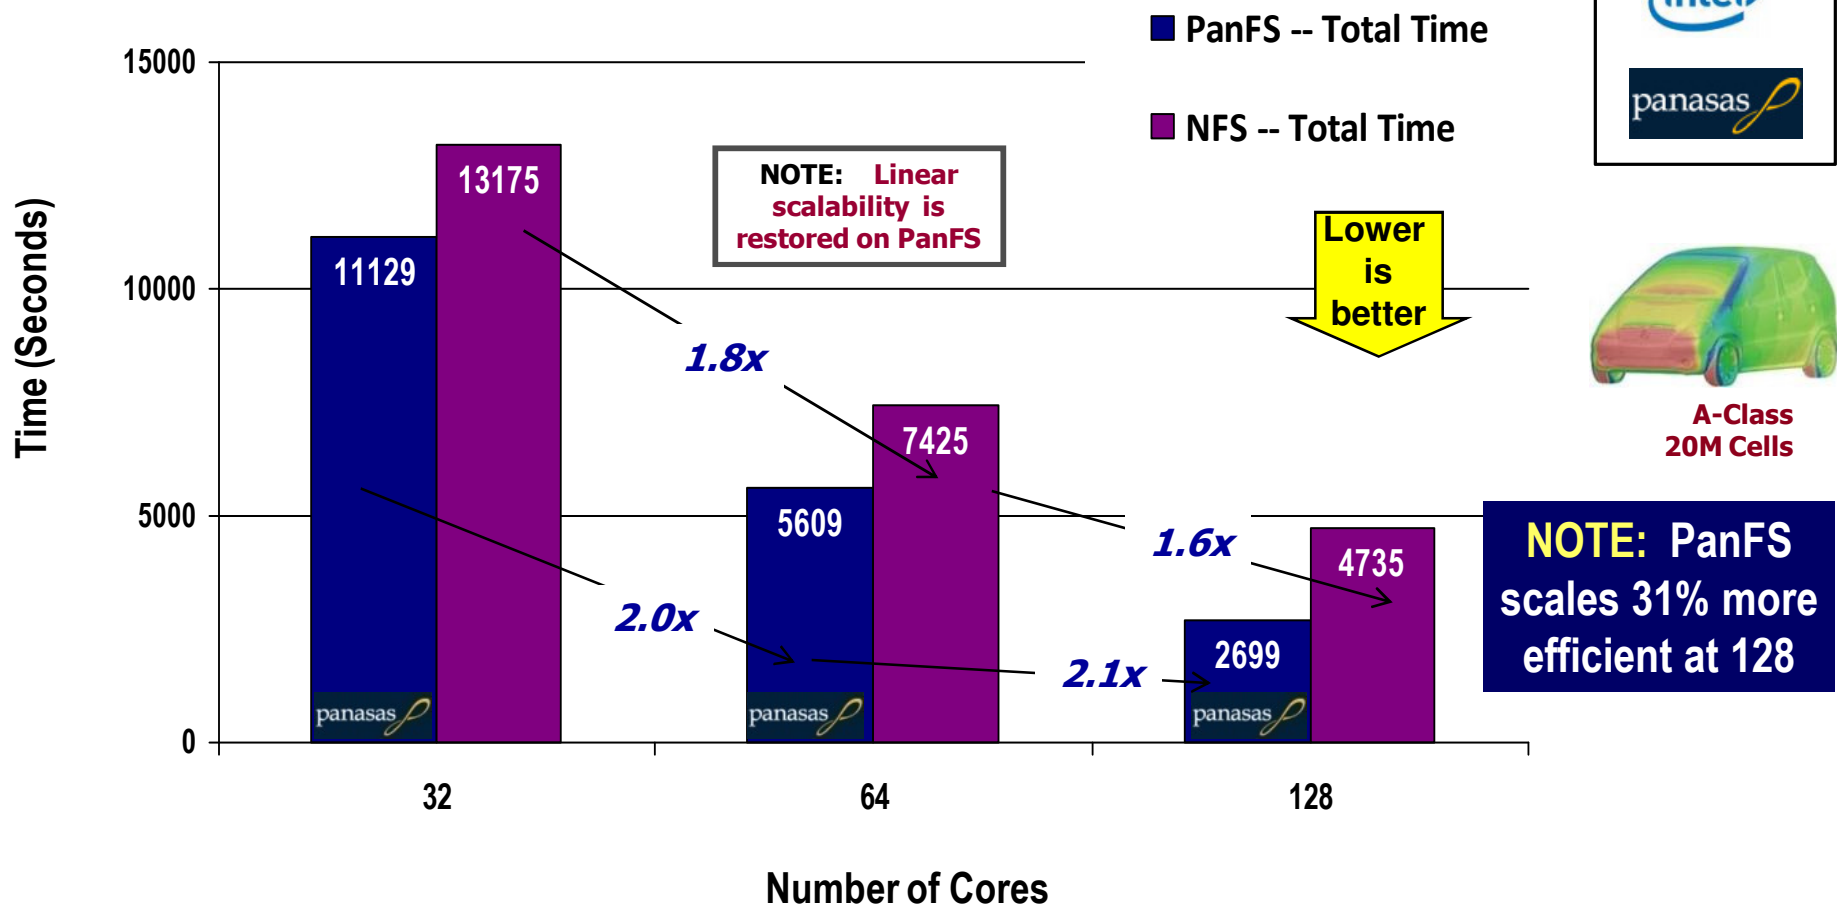
A-Class
20M Cells

NOTE: PanFS benefits grow with more cores

STAR-CD v4 Scalability of Total Times



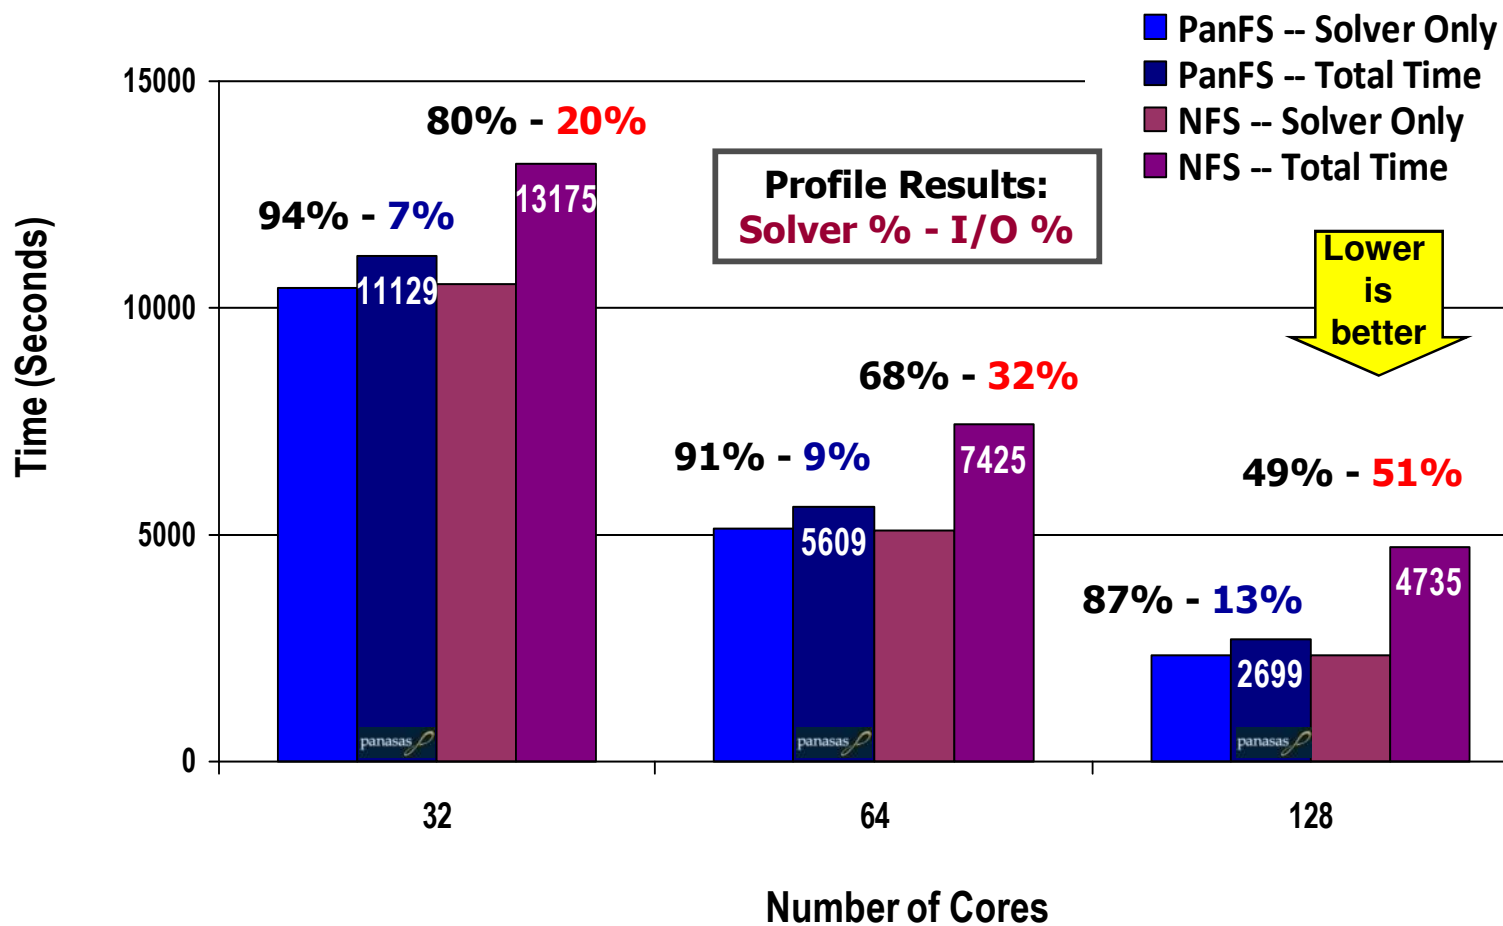
STAR-CD v4.06: Comparison of PanFS vs. NFS on Intel Cluster



STAR-CD v4 Computational Profiles



STAR-CD v4.06: Comparison of PanFS vs. NFS on Intel Cluster



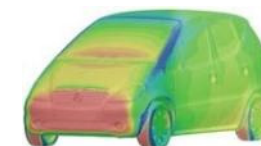
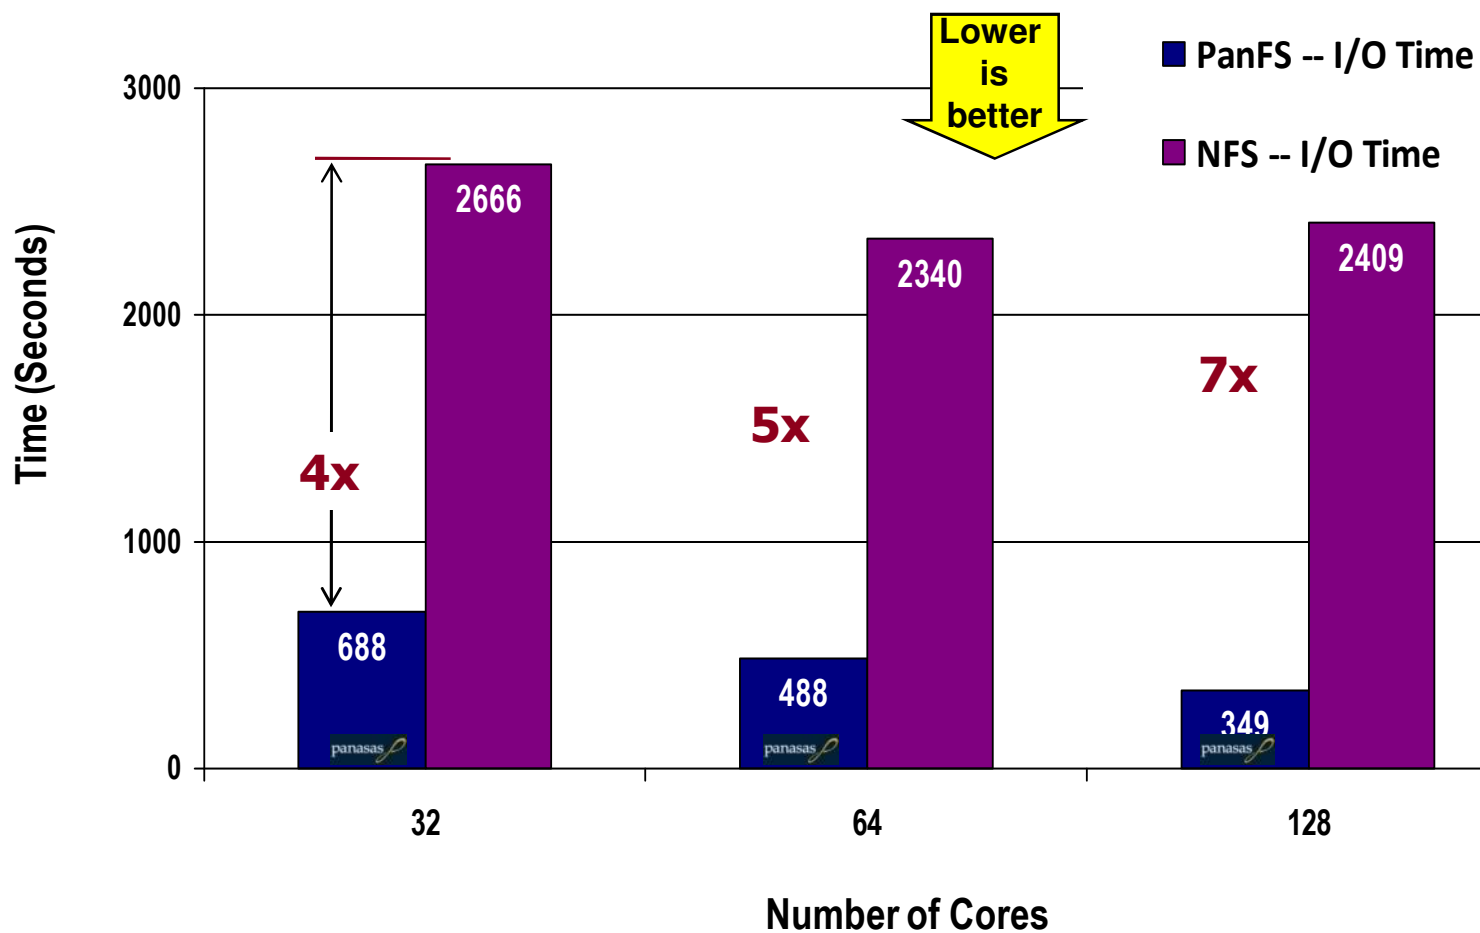
A-Class
20M Cells

NOTE: NFS profile at 128 is 51% I/O !

STAR-CD v4 Performance of I/O Times



STAR-CD v4.06: Comparison of PanFS vs. NFS on Intel Cluster



A-Class
20M Cells

NOTE: PanFS
benefits grow
with more cores

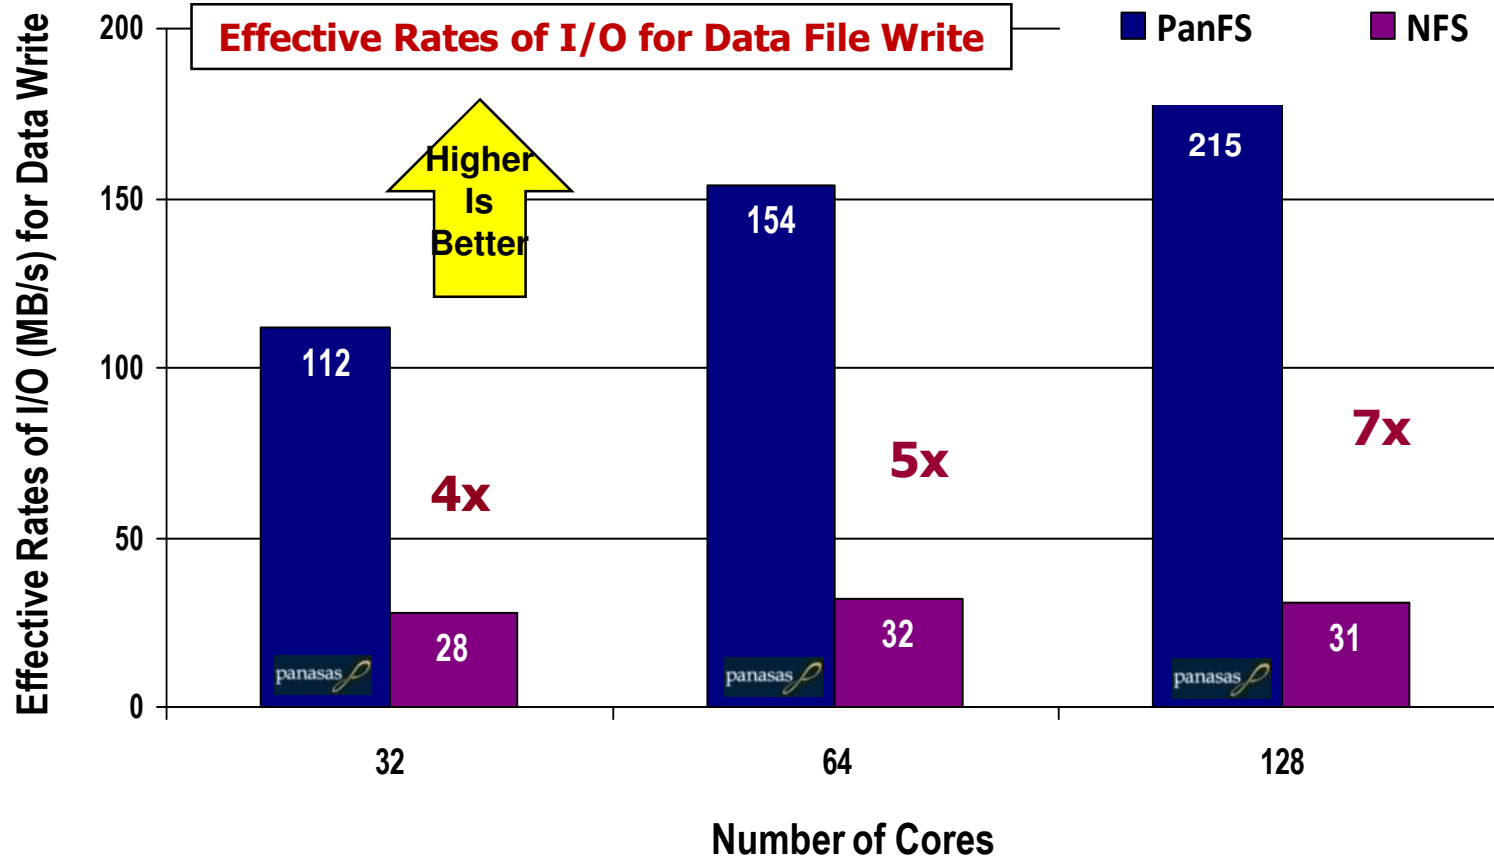
STAR-CD v4 Rates of I/O Bandwidth



STAR-CD v4.06: Comparison of PanFS vs. NFS on Intel Cluster



A-Class
20M Cells



Customer Case of STAR-CD Benchmark



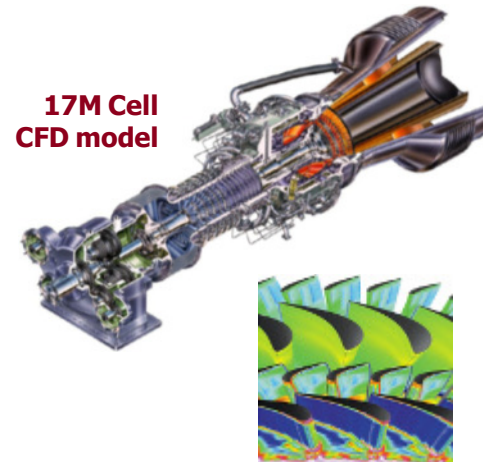
Transient solution for 17 MM cell model; 60 time steps with 300 iterations; time history writes at each 50 iterations, and solution writes at each 100 iterations

Number of cells	16,930,109
Solver	CGS, Single Precision
Iterations	300 total iterations - data save after every 100 iters
Total solution output	~48 GB

Turbomachinery Company
US-based developer of gas turbines (stationary)

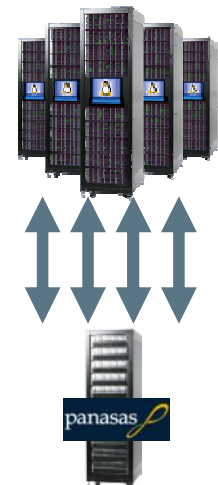


Intel "ENDEAVOR" Xeon®	
Nodes: 256 x 2 Sockets x 4 Cores = Total Cores: 2048	
Location: DuPont, WA	
CPU: Harpertown Xeon QC 2.8 GHz / 12MB L2 cache	
FSB: 1600 MHz, IB Interconnect DDR	
File System: Panasas, 7 shelves, 35 TB storage	
FS Connectivity: Gig Ethernet, 4 bonded links per shelf, 2.8 GB / sec peak, 2.5 GB./ sec measured	



17M Cell CFD model

ENDEAVOR
2048 cores

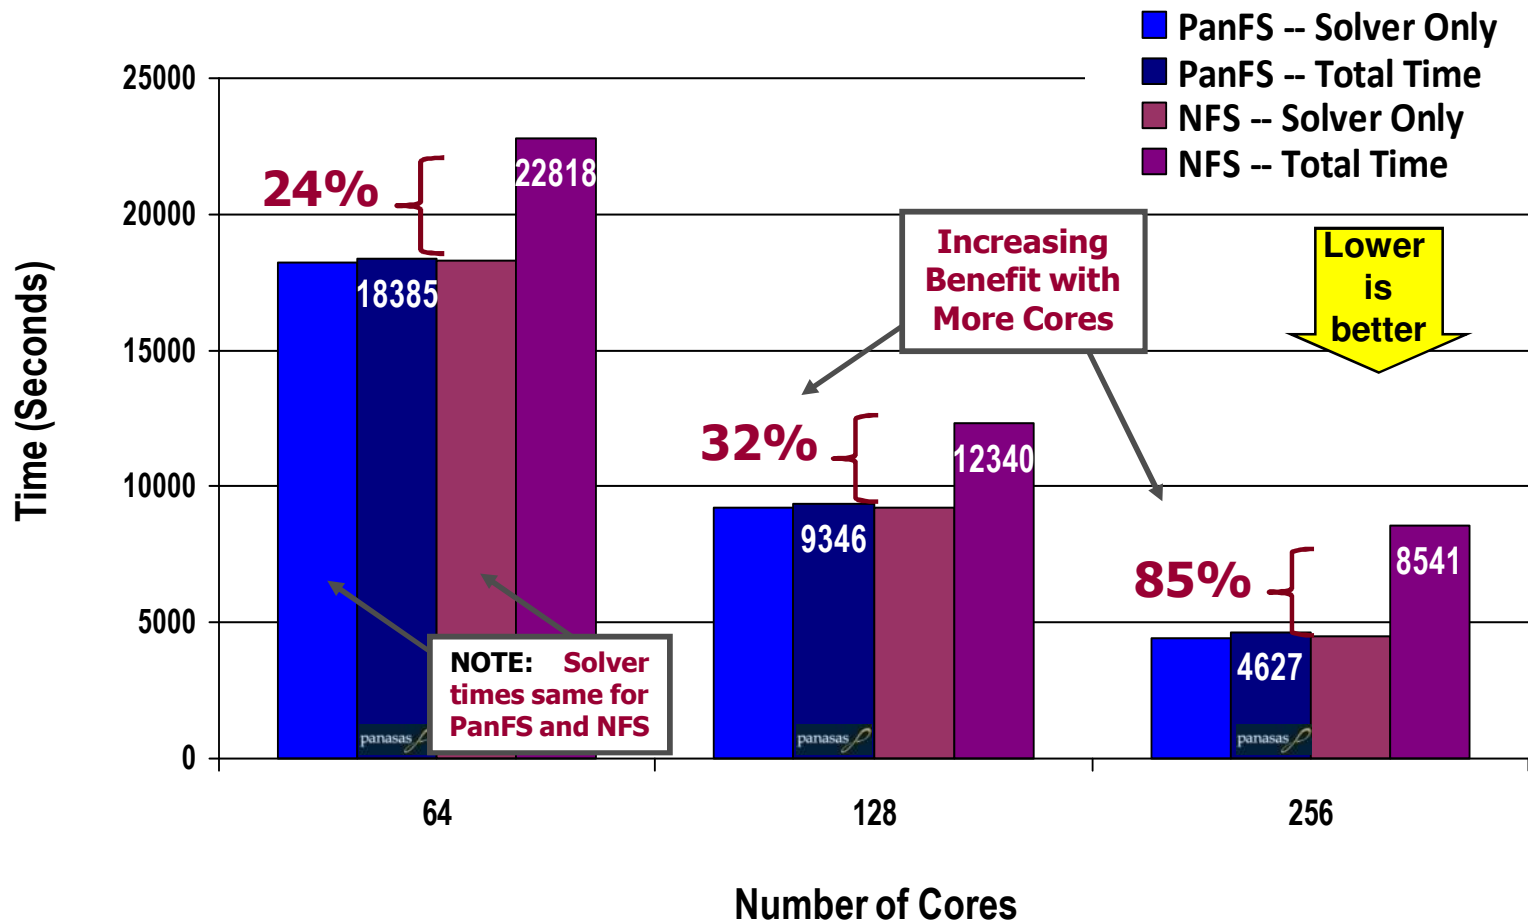


Panasas Storage:
7 Shelves, 35 TB

Customer Case of STAR-CD Benchmark



STAR-CD v4.06: Comparison of PanFS vs. NFS on Intel Clusters

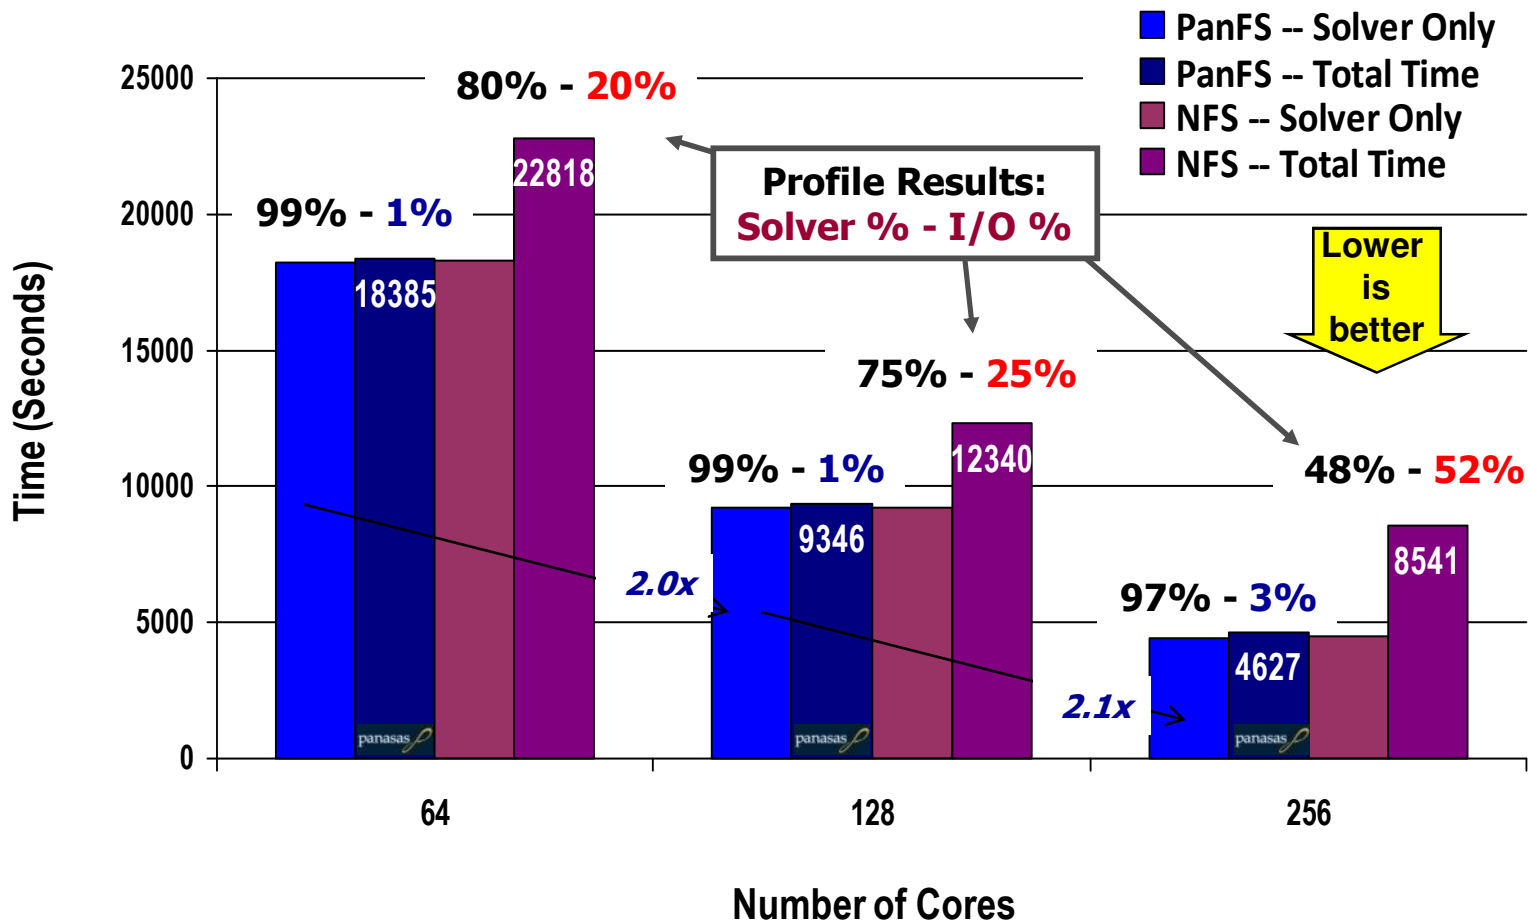


17M Cell CFD model

Customer Case of STAR-CD Benchmark



STAR-CD v4.06: Comparison of PanFS vs. NFS on Intel Clusters



- NFS and Serial I/O Limitations
 - Certain STAR-CD production cases can **waste** a substantial percentage of a computational profile in I/O operations vs. valued FP operations
 - The use of frequent checkpoints for very large steady-state cases, and/or large unsteady simulations (multiple writes) is impractical with serial I/O
- STAR-CD v4 and Panasas Solution
 - The Panasas parallel file system and storage, combined with parallel I/O of STAR-CD scales I/O and therefore the overall STAR-CD simulation
 - Use of Panasas file system for 20M cell case at 128-way provides a 75% increase in STAR-CD utilization for the same software license \$'s spent
 - Such capability enables STAR-CD users to develop more advanced CFD models (more transient vs. steady, LES, etc.) with confidence in scalability

Contributors to the Study



CD-adapco

- **Dr. Boris Kaludercic, Technical Staff, Parallel Development**
- **Mr. Ron Gray, Technical Staff, Benchmark Support**
- **Mr. Steve Feldman, VP, Software Development and IT**



Intel

- **Mr. Paul Work, Manager, Engineering Operations**
- **Dr. Paresh Pattani, Director of Applications Engineering**



Panasas

- **Mr. Derek Burke, Director of EMEA Marketing**



CD-adapco Services is a Panasas Customer

CD-adapco

Engineering Services, Plymouth MI



■ CAE Software

- CFD - STAR-CD, STAR-CCM+; CSM – Abaqus

■ HPC Solution

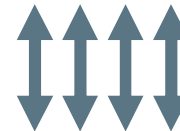
- Linux cluster (~256 cores); PanFS 30TB file system

■ Business Value

- File reads and merge operations 2x faster than NAS
- Throughput of multi-job access to the shared file system
- Parallel I/O in STAR-CD 3.26 and 4.06 can leverage the PanFS parallel file system today
- **STAR-CD v4.06 with parallel I/O released May 08**
- STAR-CCM+ has plans for parallel I/O *sim* file writes



Linux x86_64, 256 cores



Panasas: 2 Shelves, 20 TB

Panasas and CD-adapco Driving Parallel I/O

Panasas and CD-adapco Press Release – 10 Mar 2008



panasas  MYPANASAS TAKE A PRODUCT TOUR CONTACT PANASAS

Accelerating Time to Results™

HOME PRODUCTS SOLUTIONS PARTNERS SERVICES COMPANY Search

Company Overview
Management
News & Events
▶ Press Releases
▶ Press Kit
▶ Recent Articles
▶ Upcoming Events
▶ Webinars
Awards

Company > News & Events > Press Release

Panasas and CD-adapco Partner to Advance CAE Productivity

Panasas Parallel Storage Heralds Breakthrough Performance with CD-adapco Software

FREMONT, Calif.—March 10, 2008—Panasas, Inc., the global leader in parallel storage solutions for the High Performance Computing market, announced today that it has forged a strategic alliance with Computer Aided Engineering (CAE) market leader CD-adapco. Recent certification of CD-adapco's STAR-CD and STAR-CCM+ Computational Fluid Dynamics (CFD) software on Panasas® ActiveStor parallel storage delivers new and significant performance advantages for a broad range of CAE simulations. The industry benefit is faster time to solution, which allows companies and organizations within the automotive, aerospace, turbo machinery, oil and gas, and other industries to be more productive and more profitable.



[Back to Press Release](#)

"We are delighted with our Panasas collaboration as it delivers immediate improvements in simulation scalability and workflow efficiency for our customers," said [Steve MacDonald](#), president and co-founder of CD-adapco. "The performance advantages of a Panasas and CD-adapco solution validate our shared commitment to addressing the most demanding CAE simulation requirements. This solution has helped us meet the expanding CAE objectives of our customers."

Source: www.panasas.com

Panasas Industry Leadership in HPC



US DOE: Panasas selected for *Roadrunner*, ~2PB file system – top of Top 500

- LANL \$133M system for weapons research: www.lanl.gov/roadrunner



SciDAC: Panasas CTO selected to lead Petascale Data Storage Institute

- CTO Gibson leads PDSI launched Sep 06, leveraging experience from PDSI members: LBNL/NERSC; LANL; ORNL; PNNL; Sandia NL; CMU; UCSC; UoMI



Aerospace: Airframes and engines, both commercial and defense

- Boeing HPC file system; 3 major engine mfg; top 3 U.S. defense contractors



Formula-1: HPC file system for Top 2 clusters – 3 teams in total

- Top clusters at an F-1 team with a UK HPC center and BMW Sauber




Intel: Certified Panasas storage for range of HPC applications – *Panasas Now ICR*

- Intel is a customer, uses Panasas storage in EDA and HPC benchmark center



SC08: Panasas won 5 of the annual HPC Wire Editor's and Reader's Choice Awards

- Awards for roadrunner (3) including "Top Supercomputing Achievement" 
- "Top 5 vendors to watch in 2009" | "Reader's Best HPC Storage Product"



Validation: Panasas customers won 8 out of 12 HPC Wire industry awards for SC08:



Boeing



Renault F1



Citadel



Ferrari F1



Fugro



NIH



PGS



WETA

Thank You for This Opportunity



RESOURCES:

- Questions can be directed to the Panasas email addresses below
- The 20M cell A-class model is public and available from CD-adapco
<http://www.cd-adapco.com/>
- STAR-CD log files of all jobs are available upon request to Panasas

Stan Posey
sposey@panasas.com

Bill Loewe
bloewe@panasas.com