

# **Panasas: High Performance Storage for the Engineering Workflow**

E. Jassaud, W. Szoecs

Panasas / transtec AG



*The leader in high performance parallel storage  
for business-critical applications*

## **High-Performance Storage for the Engineering Workflow**

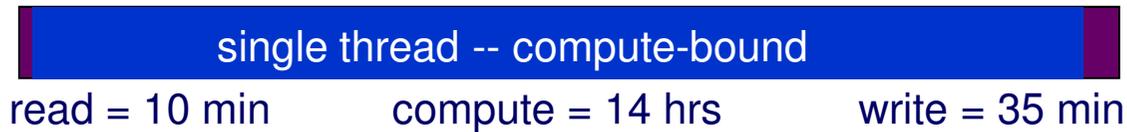
Terry Rush

Regional Manager, Northern Europe.

# Why is storage performance important?

## Example - Progression of a Single Job CAE Profile with non-parallel I/O

1998: Desktops



5% I/O

2003: SMP Servers



30% I/O

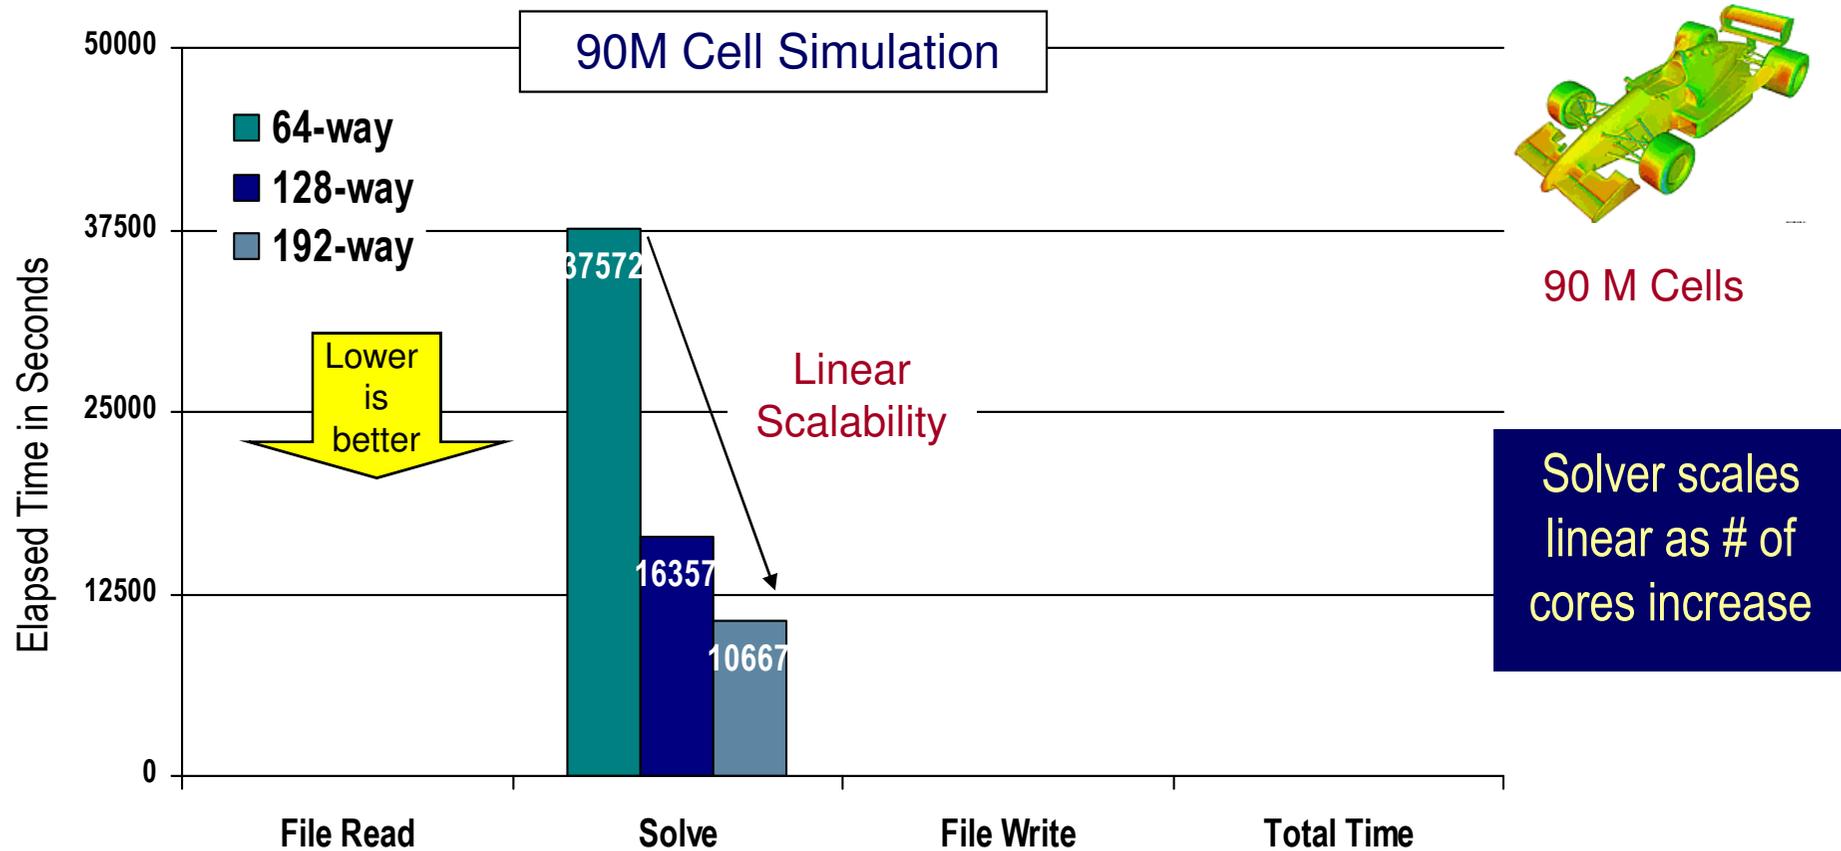
2008: HPC Clusters



63% I/O

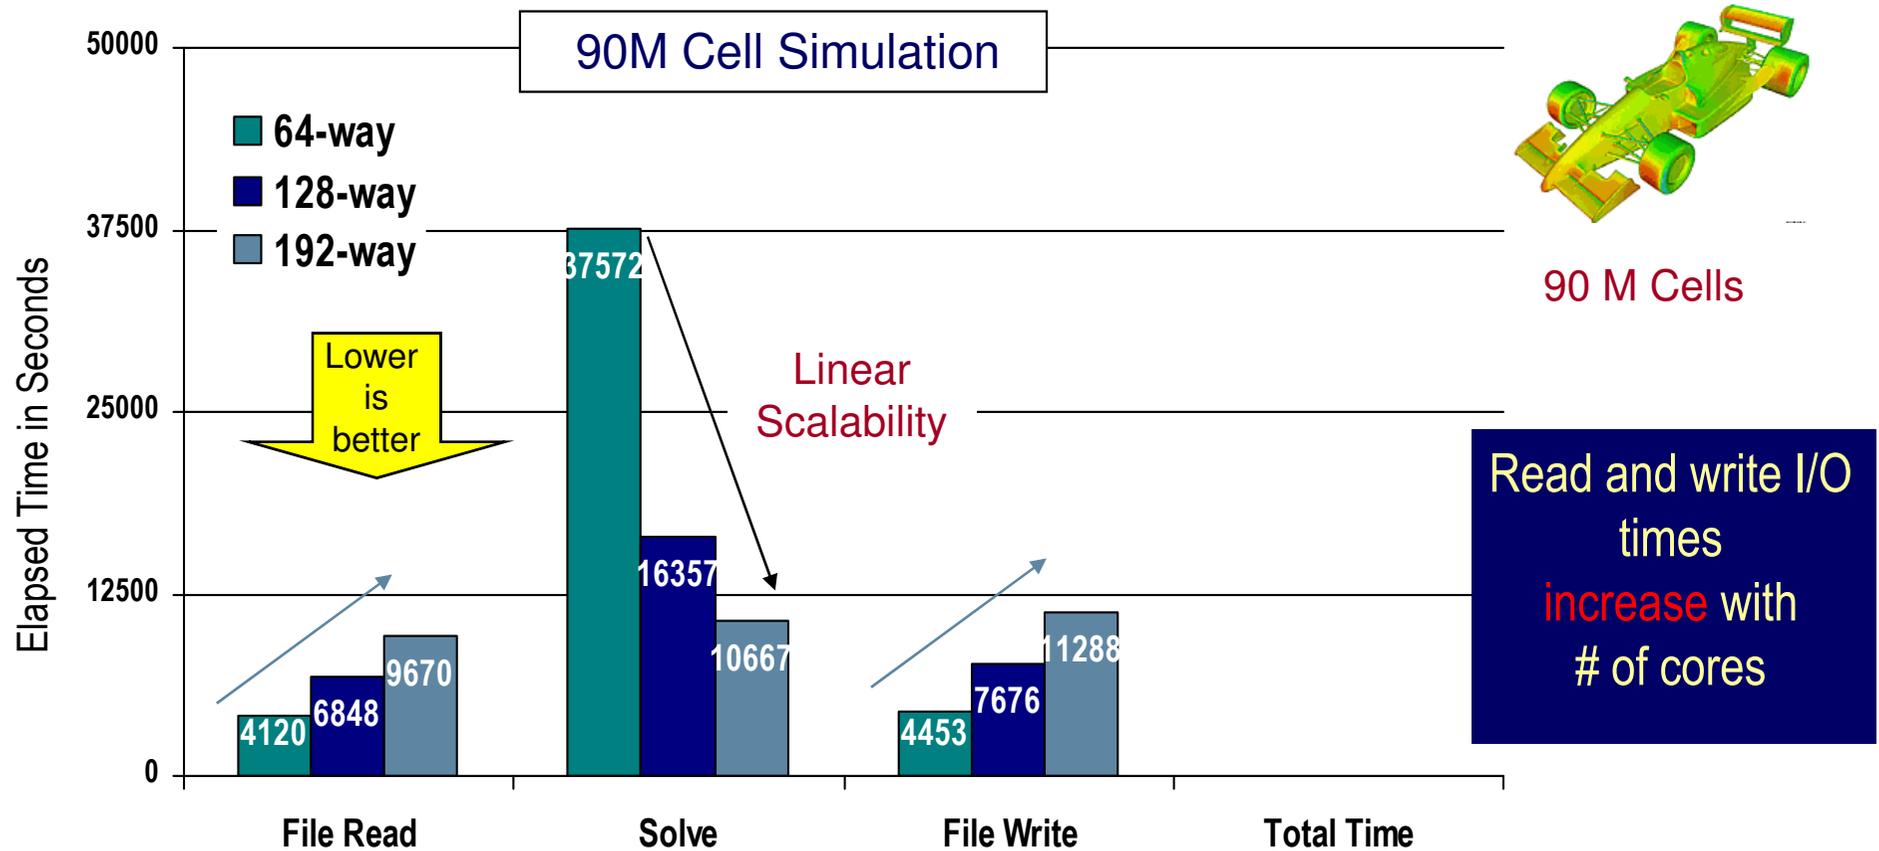
NOTE: Schematic only, hardware and CAE software advances have been made on non-parallel I/O

# The I/O problem



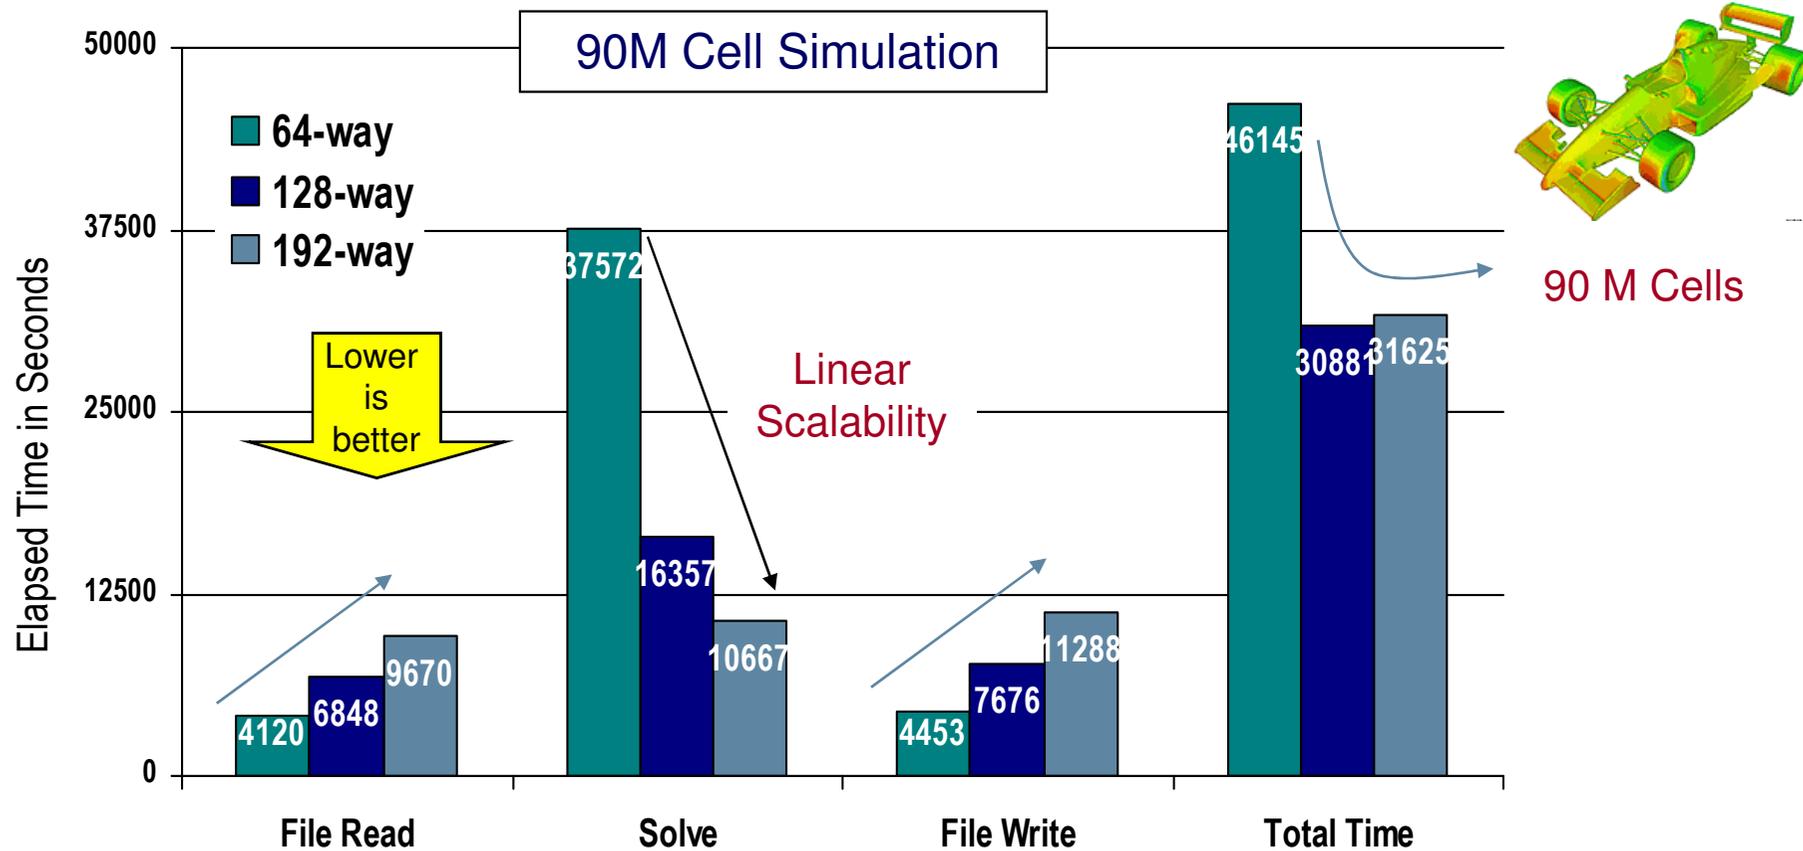
Source: Barb Hutchings Presentation at SC07, Nov 2007, Reno, NV

# Read & Write Time increases with Core Count



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

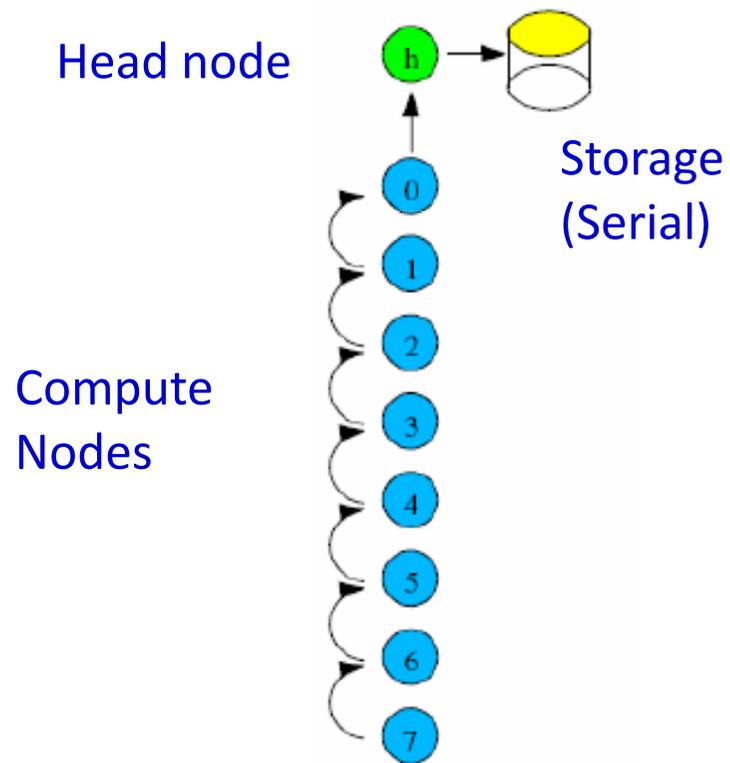
# Bigger cluster does not mean faster jobs



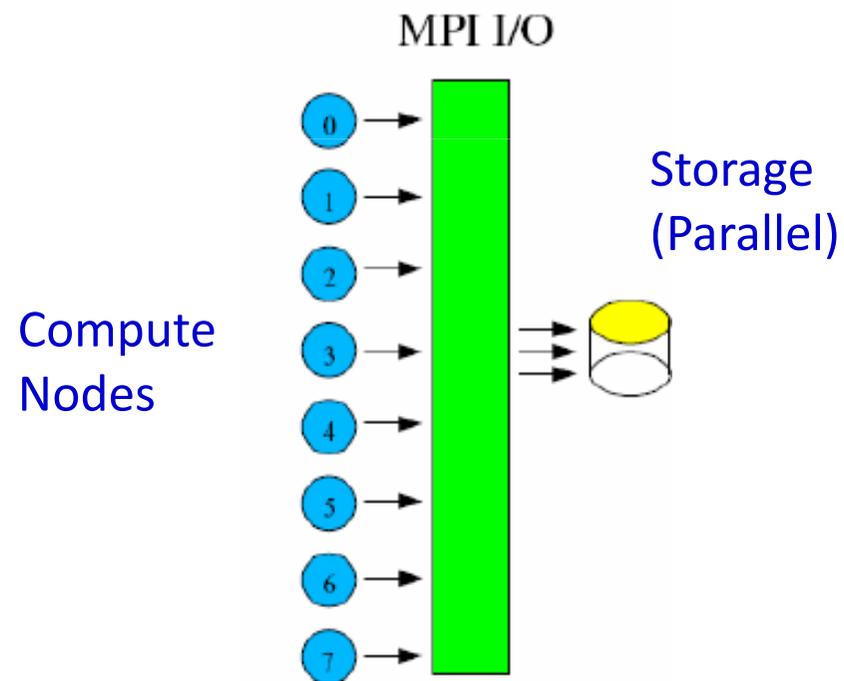
Source: Barb Hutchings Presentation at SC07, Nov 2007, Reno, NV

# Serial verses Parallel

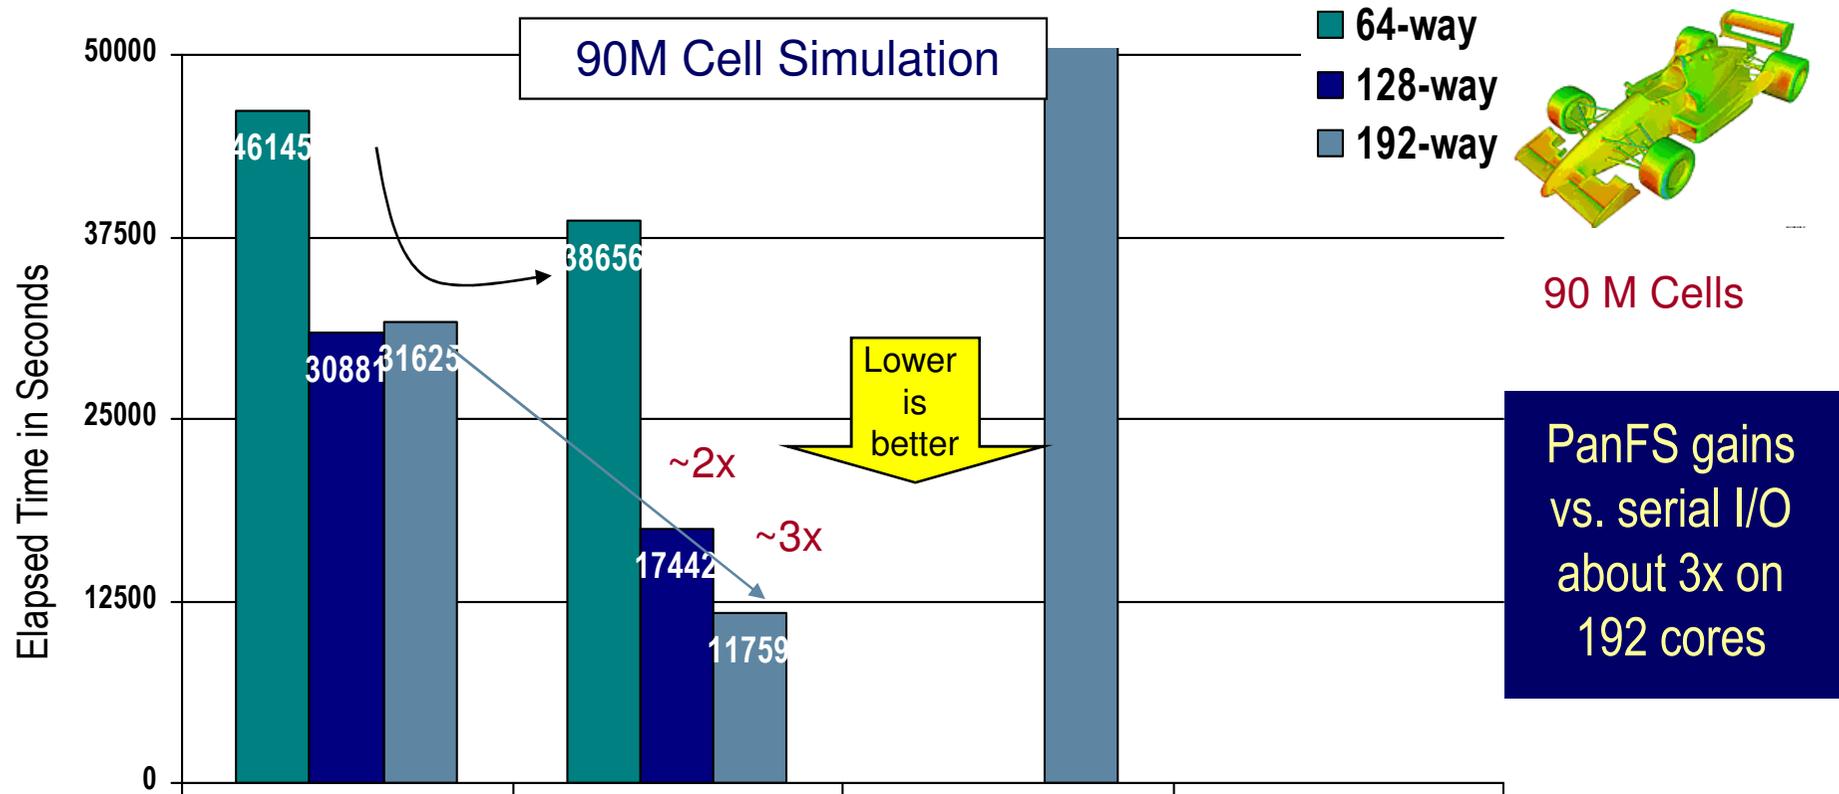
Serial I/O Schematic



Parallel I/O Schematic



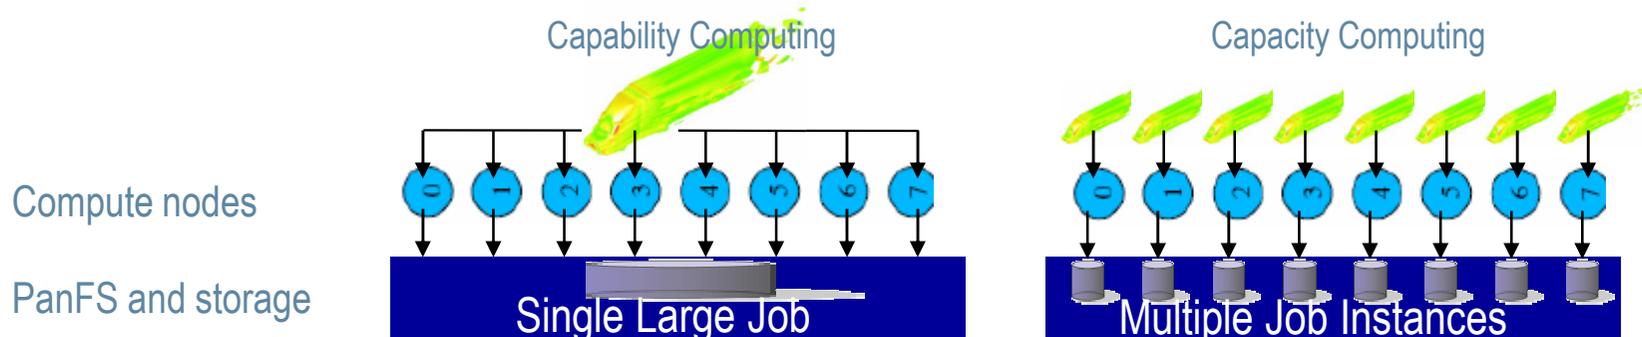
# Parallel storage I/O returns scalability



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

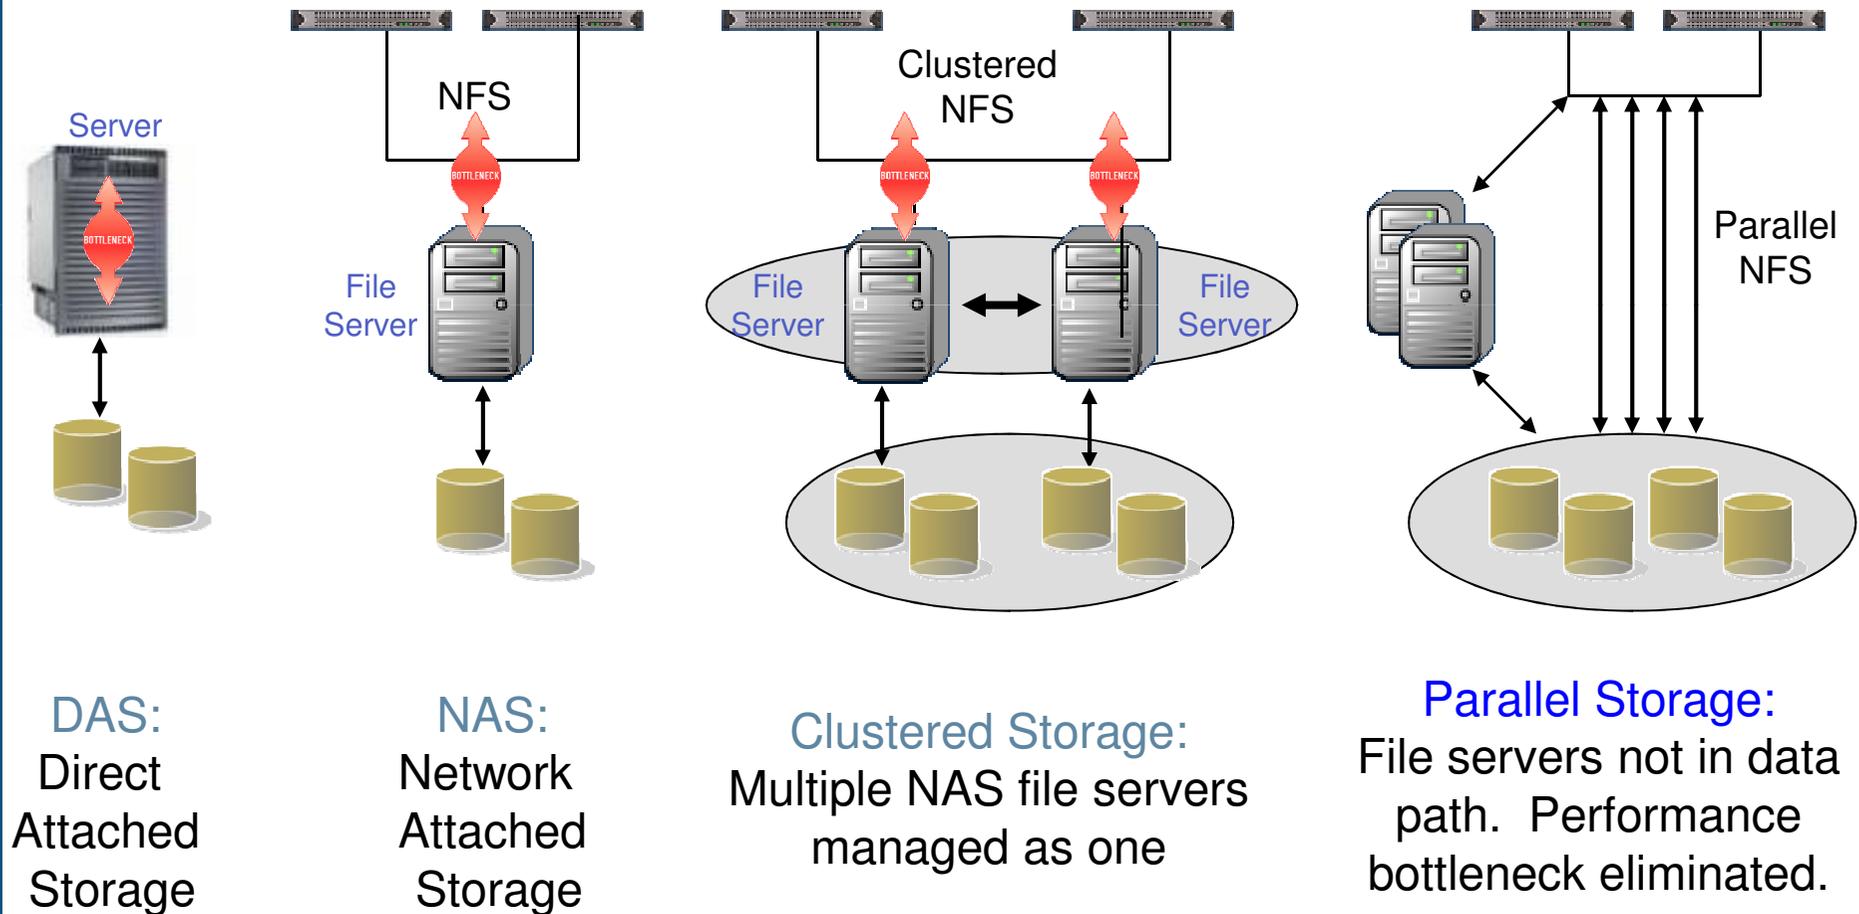
# Two Measures of Performance

- Single wide-parallel job vs. multiple instances of less-parallel jobs
  - Often referenced in HPC industry as capability vs. capacity computing
  - Both important, but multi-user capacity computing more common in practice
    - Example: parameter studies and design optimization usually based on capacity computing since it would be economically impractical with capability approach
- File system performance characteristics for capability vs. capacity
  - Capability: Many cores writing to a large single file
  - Capacity: A few cores writing into a single file, but multiple instances
  - PanFS scales I/O for the capability job, and provides multi-level scalable I/O for the capacity jobs (each job parallel I/O -- multiple jobs writing to PanFS in parallel)
  - NFS has single, serial data path for all capacity and capability solution writes



# What is Parallel Storage?

## Parallel NAS: Next Generation of HPC File System



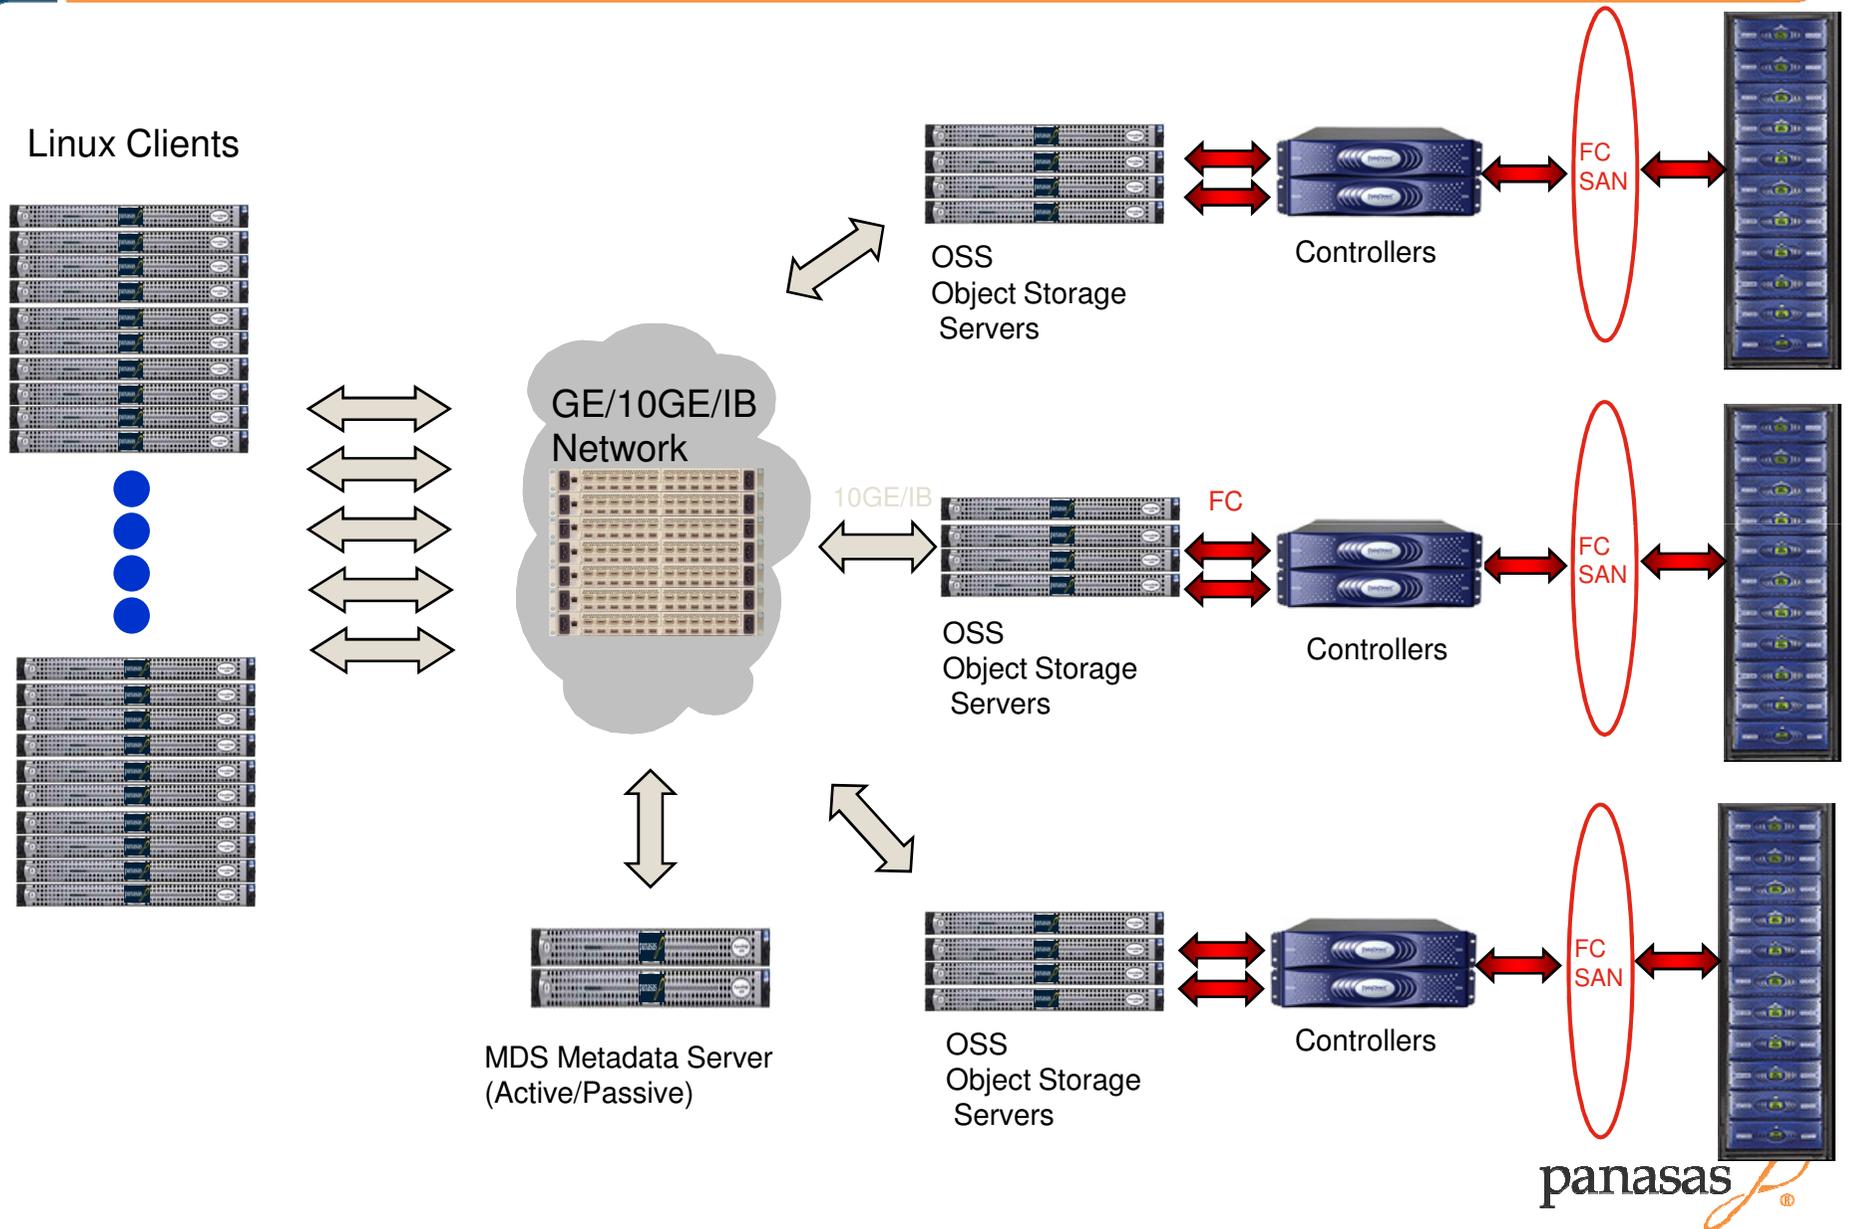
DAS:  
Direct  
Attached  
Storage

NAS:  
Network  
Attached  
Storage

Clustered Storage:  
Multiple NAS file servers  
managed as one

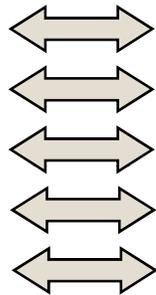
Parallel Storage:  
File servers not in data  
path. Performance  
bottleneck eliminated.

# Typical Complex Parallel Storage Configuration

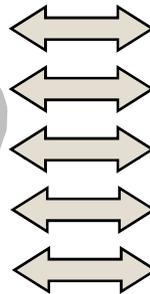
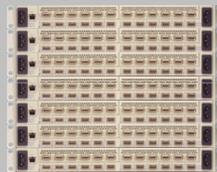


# Panasas Solution

Linux Clients



GE/10GE/IB  
Network



## Integrated solution

- Panasas File System
- Panasas HW and SW
- Panasas Service/Support
- Single management point

## Simplified Manageability

- Easy and quick to install
- Single Fabric
- Single Web GUI/CLI
- Automatic capacity Balancing

## Scalable Performance

- Scalable Metadata
- Scalable NFS/CIFS
- Scalable Reconstruction
- Scalable Clients
- Scalable BW

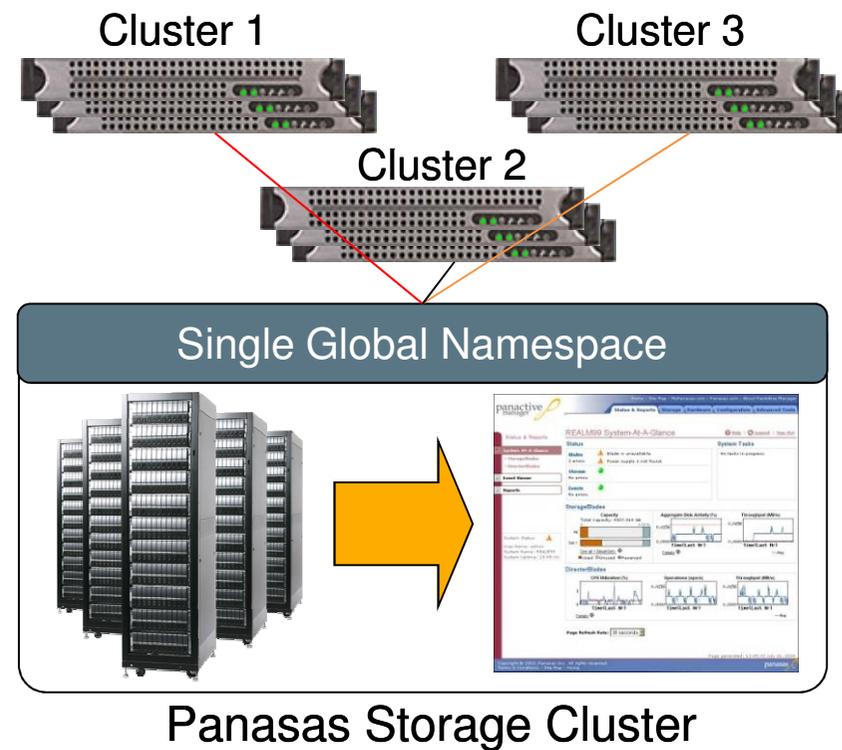
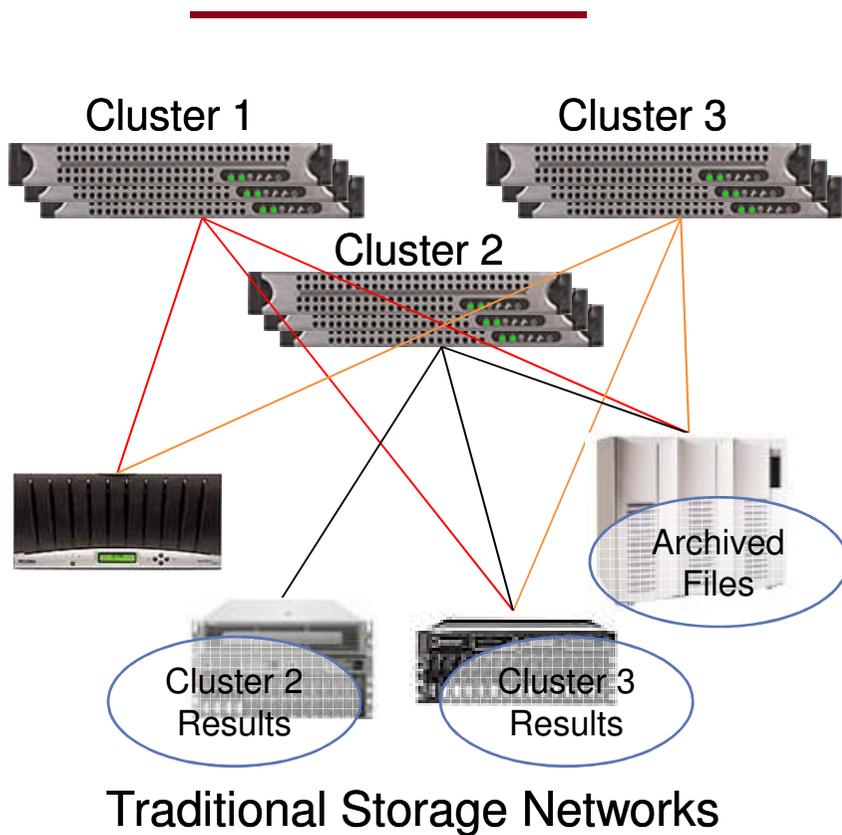
## High Availability

- Metadata failover
- NFS failover
- Snapshots
- NDMP

# What else besides fast I/O? Single Global Namespace for Easy Management

## ❖ Remove artificial, physical and logical boundaries

- Eliminates need to maintain mount scripts or move data



# Removing Silos of Storage

## ❖ Historic technical computing storage problem

- Multiple applications and users across the workflow creates costly and inefficient storage silos by application and by vendor

### Supercomputing

*HPC Linux Cluster*



High Throughput  
Sequential I/O

### High Capacity Store

*Linux/Windows/Unix*



Low \$/GB  
Massive Capacity

### Technical Apps

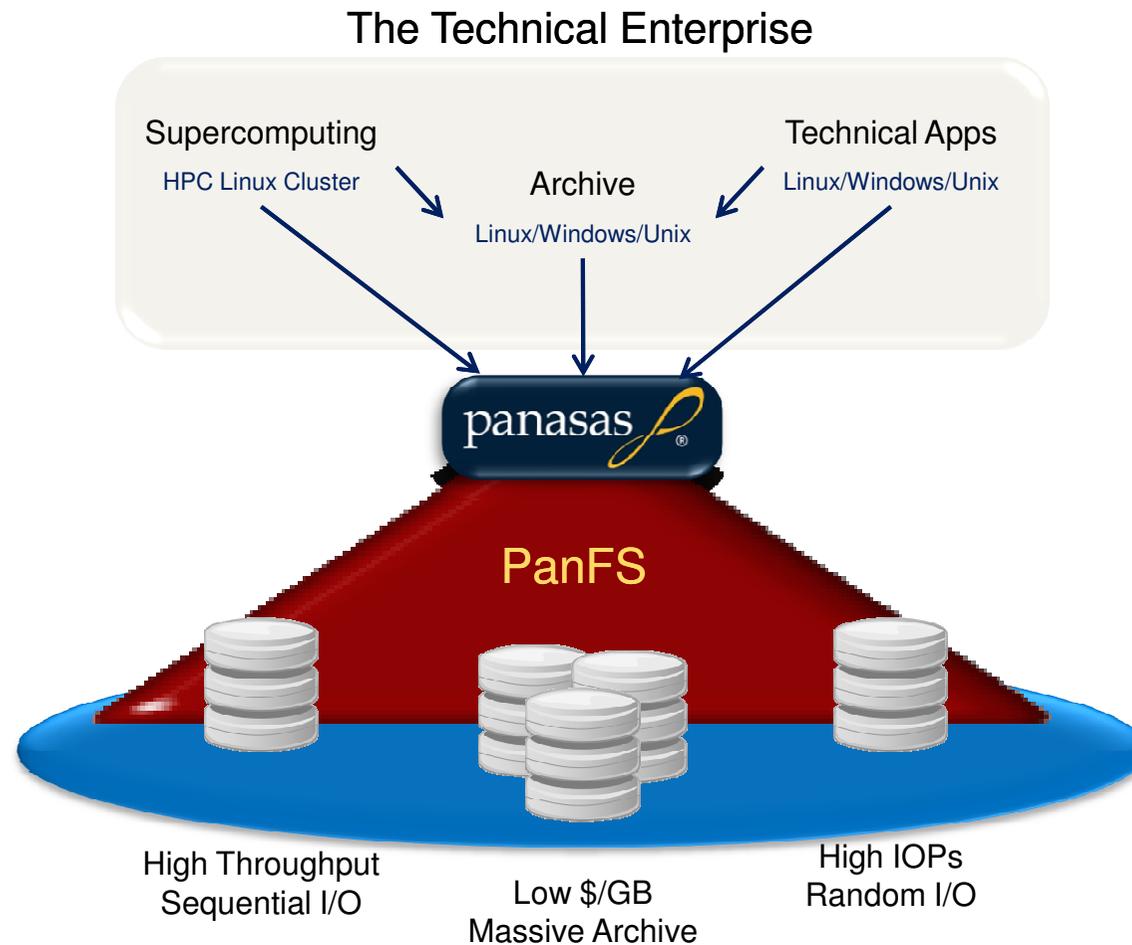
*Linux/Windows/Unix*



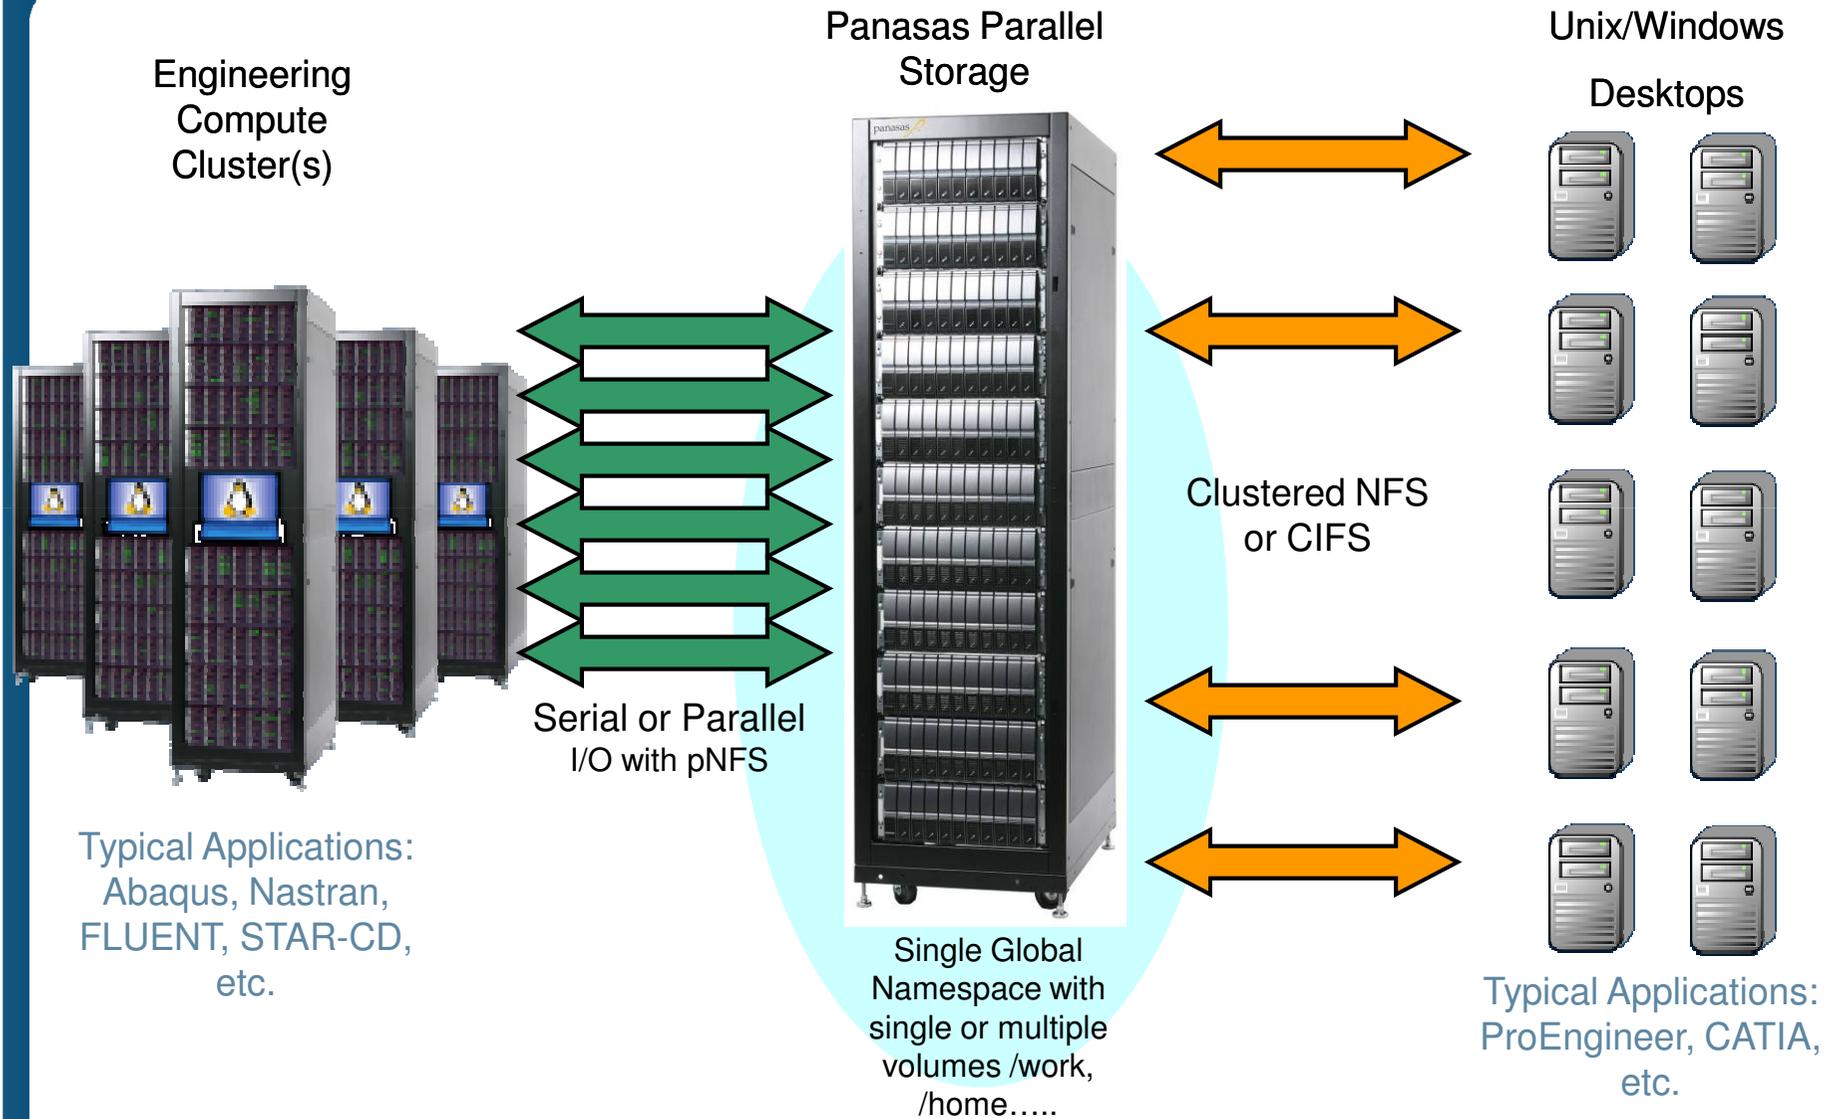
High IOPs  
Random I/O

# Removing Silos of Storage – single global file system

- High performance scale-out NAS solutions
  - PanFS Storage Operating System delivers exceptional performance, scalability & manageability

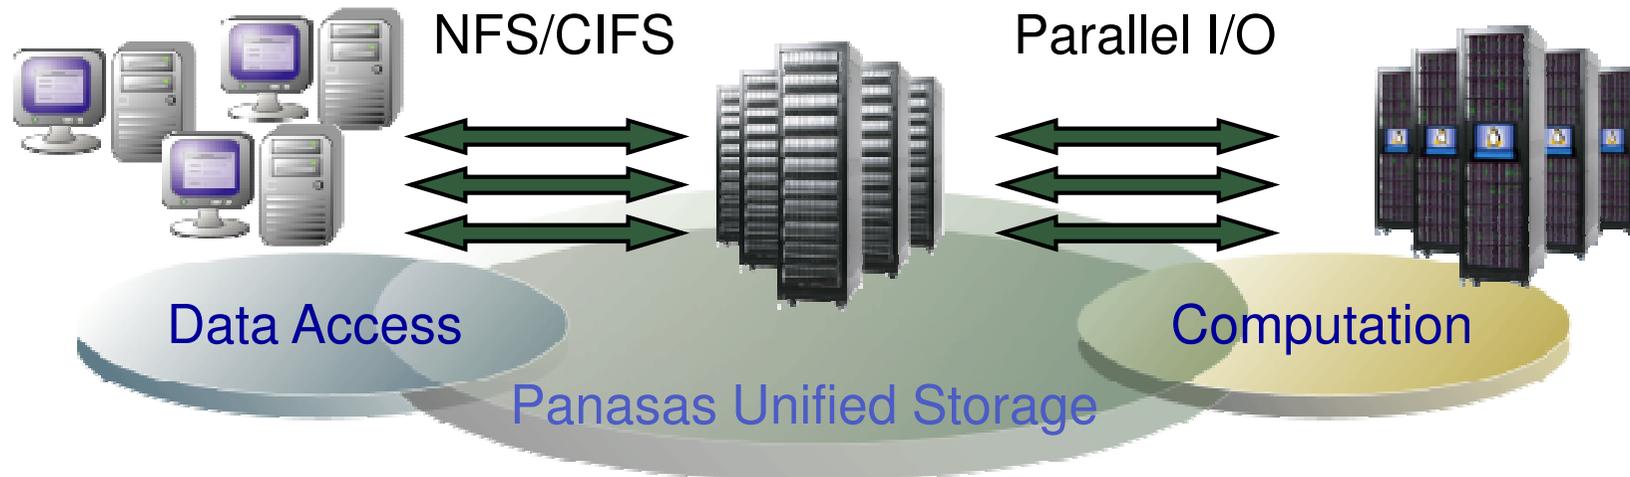


# Typical Panasas Parallel Storage Workflow:



# Unified Parallel Storage for the Engineering Workflow

- Panasas Unified Parallel Storage for Leverage of HPC investments
  - Performance that meets growing I/O demands
  - Management simplicity and configuration flexibility for all workloads
  - Scalability for fast and easy non-disruptive growth
  - Reliability and data integrity – for long term storage



# Panasas Overview

- Founded April 1999 by Prof. Garth Gibson, CTO
- Locations
  - US: HQ in Fremont, CA  
Development center in Pittsburgh, PA
  - EMEA: Sales offices in UK, France & Germany
  - APAC: Sales offices in China & Japan
- First Customer Shipments October 2003, now serving over 400 customer sites with over 500,000 server clients in twenty-seven countries
- 13 patents issued, others pending

- Key Investors



- Speed with Simplicity and Reliability - Parallel File System and Parallel Storage Appliance Solutions



# Broad Industry and Business-Critical Application Support



## Energy

Seismic Processing  
Reservoir Simulation  
Interpretation



## Government

Imaging & Search  
Stockpile Stewardship  
Weather Forecasting



## Universities

High Energy Physics  
Molecular Dynamics  
Quantum Mechanics



## Aerospace

Fluid Dynamics (CFD)  
Structural Mechanics  
Finite Element Analysis



## Finance

Credit Analysis  
Risk Analysis  
Portfolio Optimization



## Industrial Mfg

EDA Simulation  
Optical Correction  
Thermal Mechanics



## Automotive

Crash Analysis  
Fluid Dynamics  
Acoustic Analysis



## Bio/Pharma

BioInformatics  
Computation Chemistry  
Molecular Modeling

## Customer Success Highlights

- ❖ Almost half the Formula One teams use Panasas
- ❖ Five of the top six Oil & Gas companies in the world use Panasas
- ❖ Every Intel chip since 2006 has been designed using Panasas
- ❖ The world's first Petascale system, RoadRunner, uses Panasas
- ❖ The world's three largest genomic data centres use Panasas
- ❖ The largest aircraft manufacturer in the world uses Panasas
- ❖ Leading Universities including Cambridge, Oxford, Stanford & Yale use Panasas
- ❖ The world's largest Hedge Fund uses Panasas for risk analysis



Energy



Automotive



Aerospace



Bio Science  
Pharma



Finance



University  
Research



Government

# Some Panasas Customers



# Panasas Differentiation

- Leader in high performance parallel storage for business-critical applications
  - Optimized for demanding storage applications in the energy, government, finance, manufacturing, bioscience and core research sectors
- Panasas® PanFS™
  - Parallel file system with integrated RAID protection
  - Most scalable performance in the industry
  - Appliance-like scalable hardware architecture for large-scale storage deployments
  - Highly integrated management suite

# Integrated RAID Reliability

RAID Implementation	Benefit
Object RAID	<ul style="list-style-type: none"><li>✓ System intelligently assigns RAID levels based on file size</li><li>✓ Automatic transitioning from RAID 1 to RAID 5 without re-striping</li></ul>
Per File RAID	<ul style="list-style-type: none"><li>✓ High performance file reconstruction (vs. drive sector reconstruction)</li><li>✓ Rebuild in hours</li><li>✓ Parallel rebuild – all blade sets participate in RAID rebuilds</li></ul>
Horizontal (Blade) and Vertical (Disk) Parity RAID	<ul style="list-style-type: none"><li>✓ RAID within the individual drive as well as across drives</li><li>✓ Improves internal ECC capabilities</li><li>✓ Predictatively solves media errors</li><li>✓ Significantly lowers drive failure probability</li></ul>

- ❖ Only Panasas includes RAID data protection as a component of its file system

# Appliance-Like, Modular Hardware Design

Add blades, chassis, or entire racks, without system disruption, to extend performance and capacity

DirectorBlade



- CPU, cache, network
- Manages system activity
- Clustered metadata services

StorageBlade



- CPU, cache, data storage
- Enables parallel reads/writes
- Advanced caching algorithms

11-slot, 4U Chassis



- Capacity expansion from 40TB to 4PB
- Aggregate performance scales from 1.5GB/s to 150GB/s, the highest performance available
- 10GbE & IB networks
- Additional storage seamlessly integrates
- Field replaceable
- Low TCO

40U Rack



- Up to 10 chassis
- Up to 400TB (currently)

Rear of Chassis

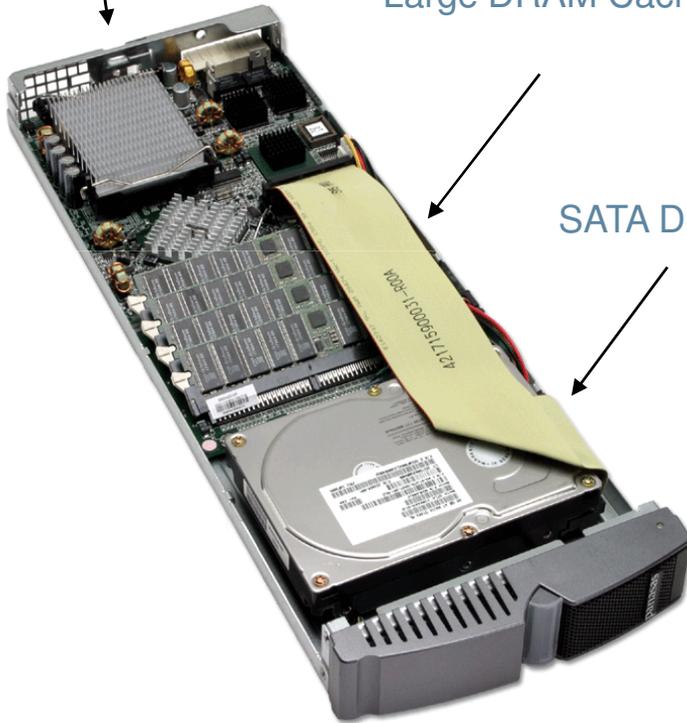


# Panasas Hardware Technology

Multi-core Intel Processor

Large DRAM Cache

SATA Drive



PAS DirectorBlade™

I/O Servers



RAID Controller



60 Drives per 4U Shelf



PAS HC Storage

# Panasas Product Family

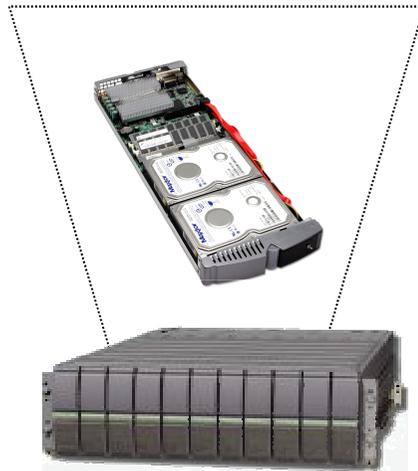
Same Panasas software, management interface, single name space



## PAS HC

(Per 42U Rack)

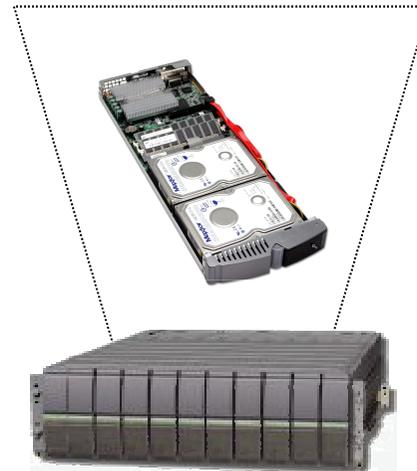
- 960TBs Raw Capacity
- 5GB/sec Aggregate B/W
- RAID 6
- Unlimited Scalability
- Entry Level \$/GB 500TBs+



## PAS 7

(Per 4U Shelf)

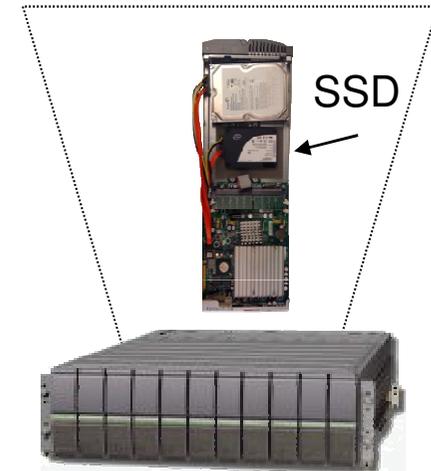
- 40TBs Raw Capacity
- 350MB/sec Aggregate B/W
- RAID 1, 5 & 10
- Unlimited Scalability
- Entry Level \$/GB Sub 500TB



## PAS 8

(Per 4U Shelf)

- 40TBs Raw Capacity
- 600MB/sec Aggregate B/W
- RAID 1, 5 & 10
- HA & Snapshots
- Unlimited Scalability

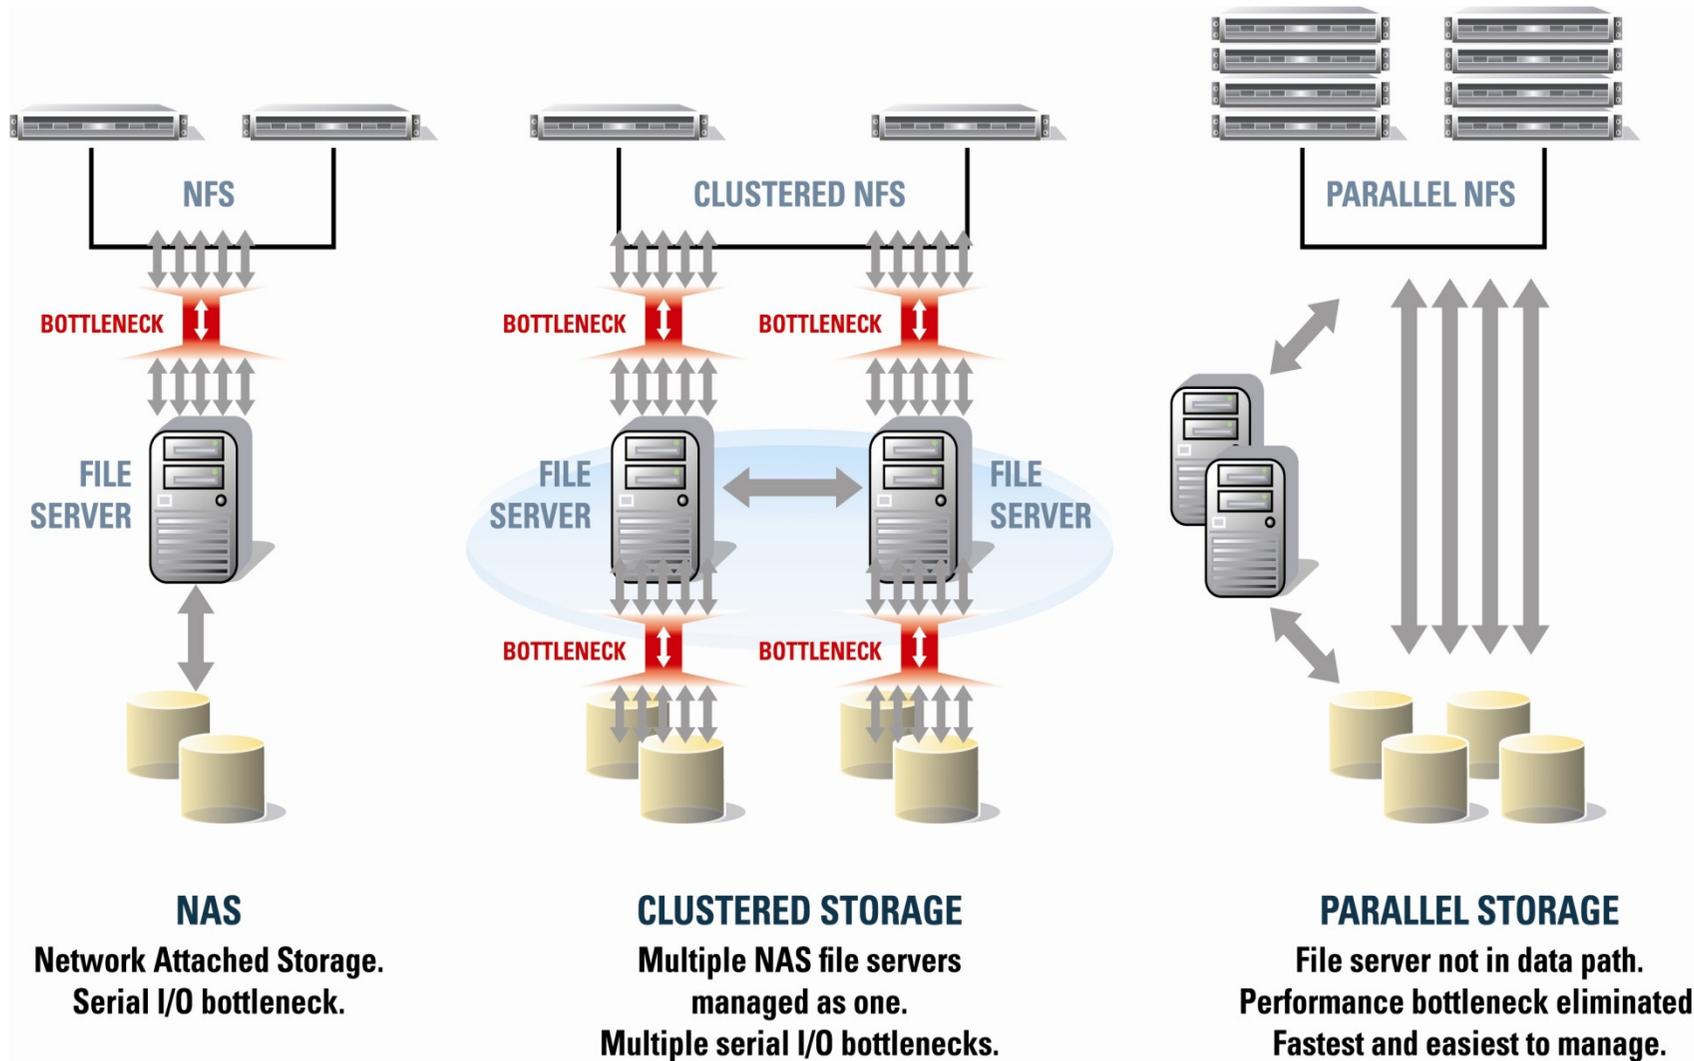


## PAS 9

(Per 4U Shelf)

- 600MB/sec Aggregate B/W
- Unlimited Scalability
- RAID 1, 5 & 10
- HA & Snapshots
- 21,000 IOPS

# High-Performance Storage for the Engineering Workflow – the Panasas Difference



Thank You!



**See us on stand 1**