Dell High Performance Cluster Computing: An Overview

Jenwei Hsieh
Dell Computer Corporation
March, 2003 @ SOS7

Product Maturity Lifecycle in the Open Systems Market

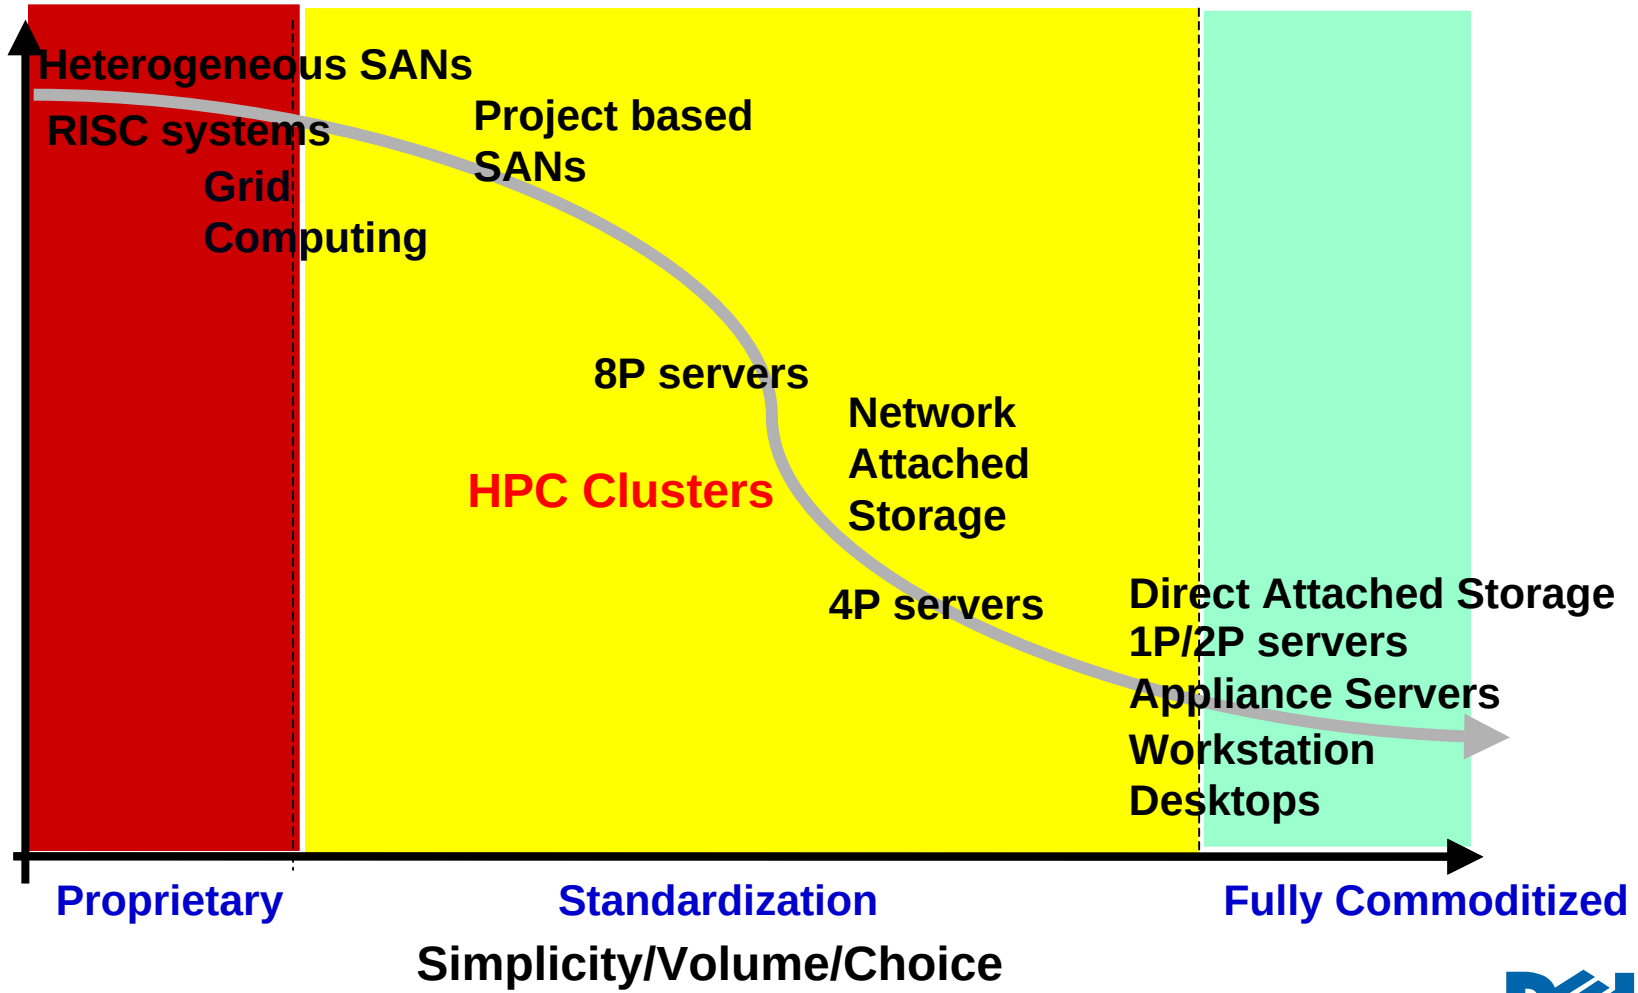

Dell HPCC Methodology

- **Baselining and Benchmarking**
- Testing Compatibility
- Tuning Performance of Components
- Developing Tools and Utilities
- Integration-Testing of Software Packages
- **Conducting R&D with Key National Labs and Universities**
- **Partnering with Best of Class Partners**
- **Sharing Our Findings**

Building Block Approach

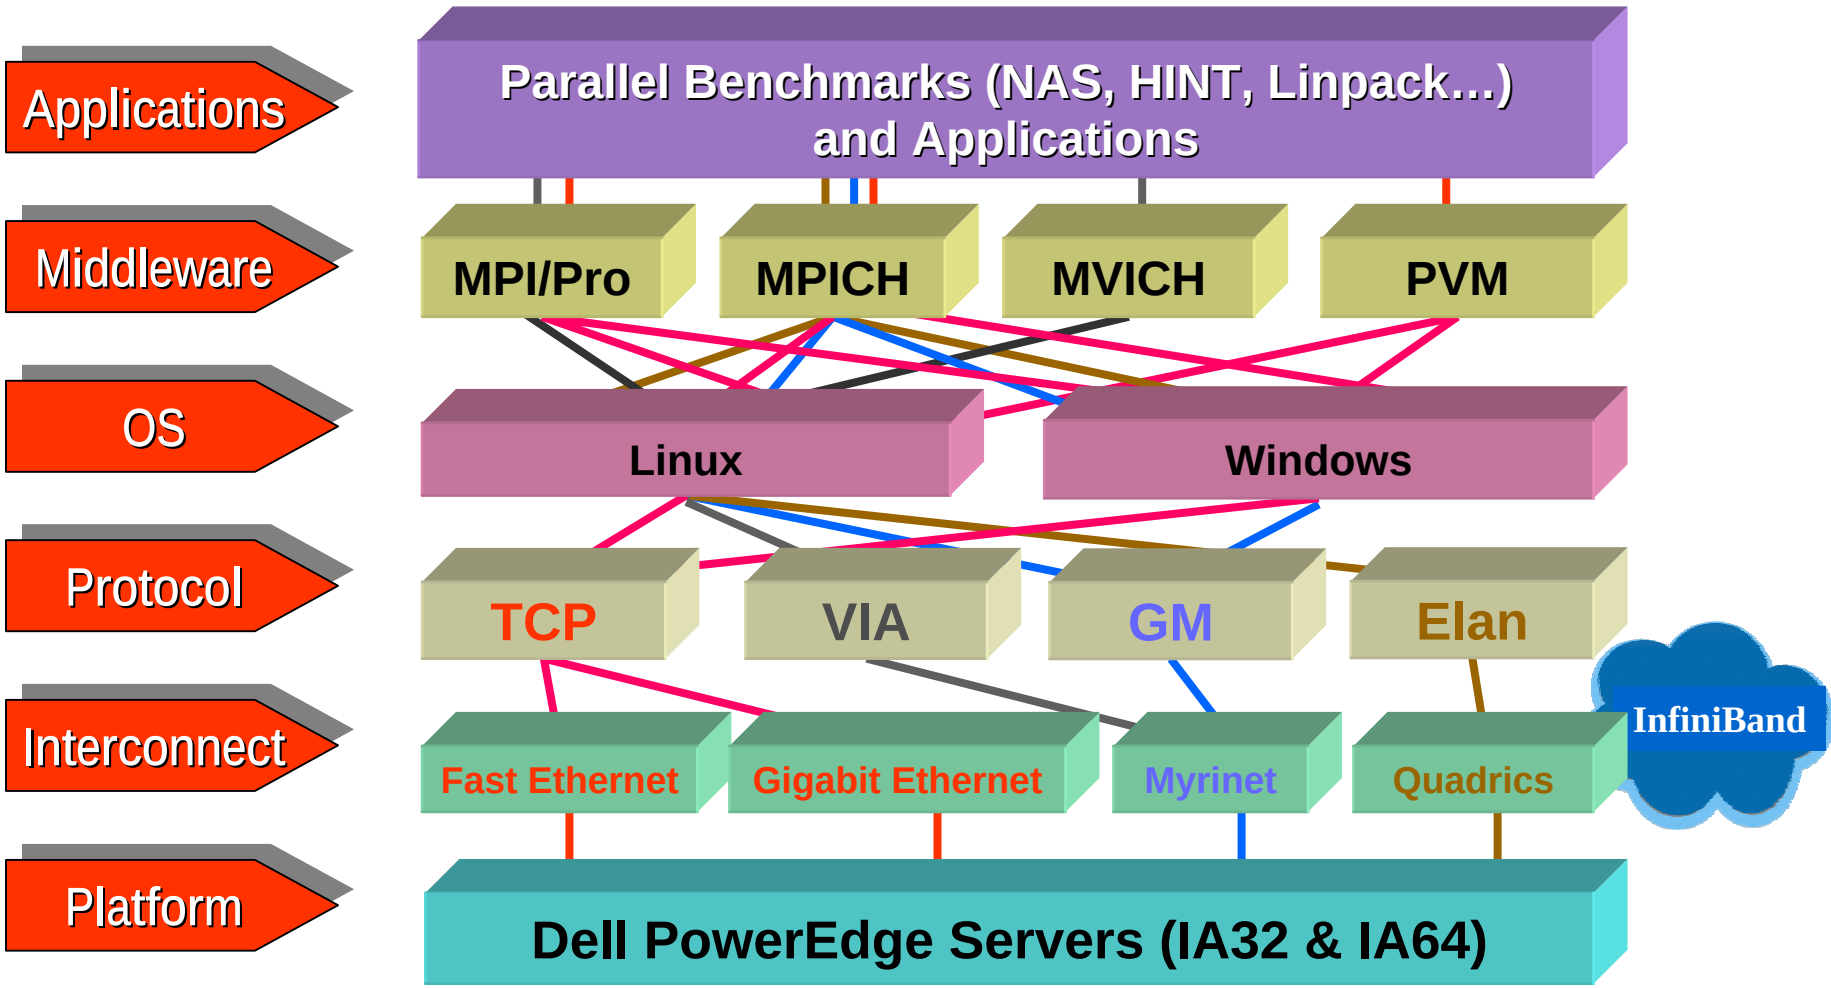

HPCC Components and Enabling Technologies

- Custom application benchmarks
- Standard benchmarks
- Performance studies

Vertical Solutions: application Prototyping / Sizing
- Energy/Petroleum - Life Science
- Automotives – Manufacturing and Design

- Reliable PVFS
- GFS , GPFS …
- Storage Cluster Solutions

Resource Monitoring / Management
Resource dynamic allocation
Checkpoint restarting and
Job redistributing

Application | Benchmark

Cluster File System

Job Scheduler

Cluster monitoring
Load analysis and Balancing
-Remote access
-Web-based GUI

Compilers and math library
Performance tools
- MPI analyzer / profiler
- Debugger
- Performance analyzer and optimizer

Development Tools

Middleware / API

Node Monitoring & Management

Cluster monitoring
Distributed System Performance Monitoring
Workload analysis and Balancing
-Remote access
-Web-based GUI

MPI 2.0 / Fault Tolerant MPI
MPICH, MPICH-GM, MPI/LAM, PVM

Operating Systems

Interconnect Protocols

Cluster Hardware Software Monitoring & Management

Interconnect Technologies
- FE, GbE, 10GE… (RDMA)
- Myrinet, Quadrics, Scali
- Infiniband

Interconnects Hardware

Management Hardware

Cluster Installation

Remote installation / configuration
PXE support
System Imager
LinuxBIOS

IA-32, IA64 (Processor / Platform) comparison
Standard rack mounted, blade and brick servers / workstations

Platform Hardware

In-the-Box Scalability

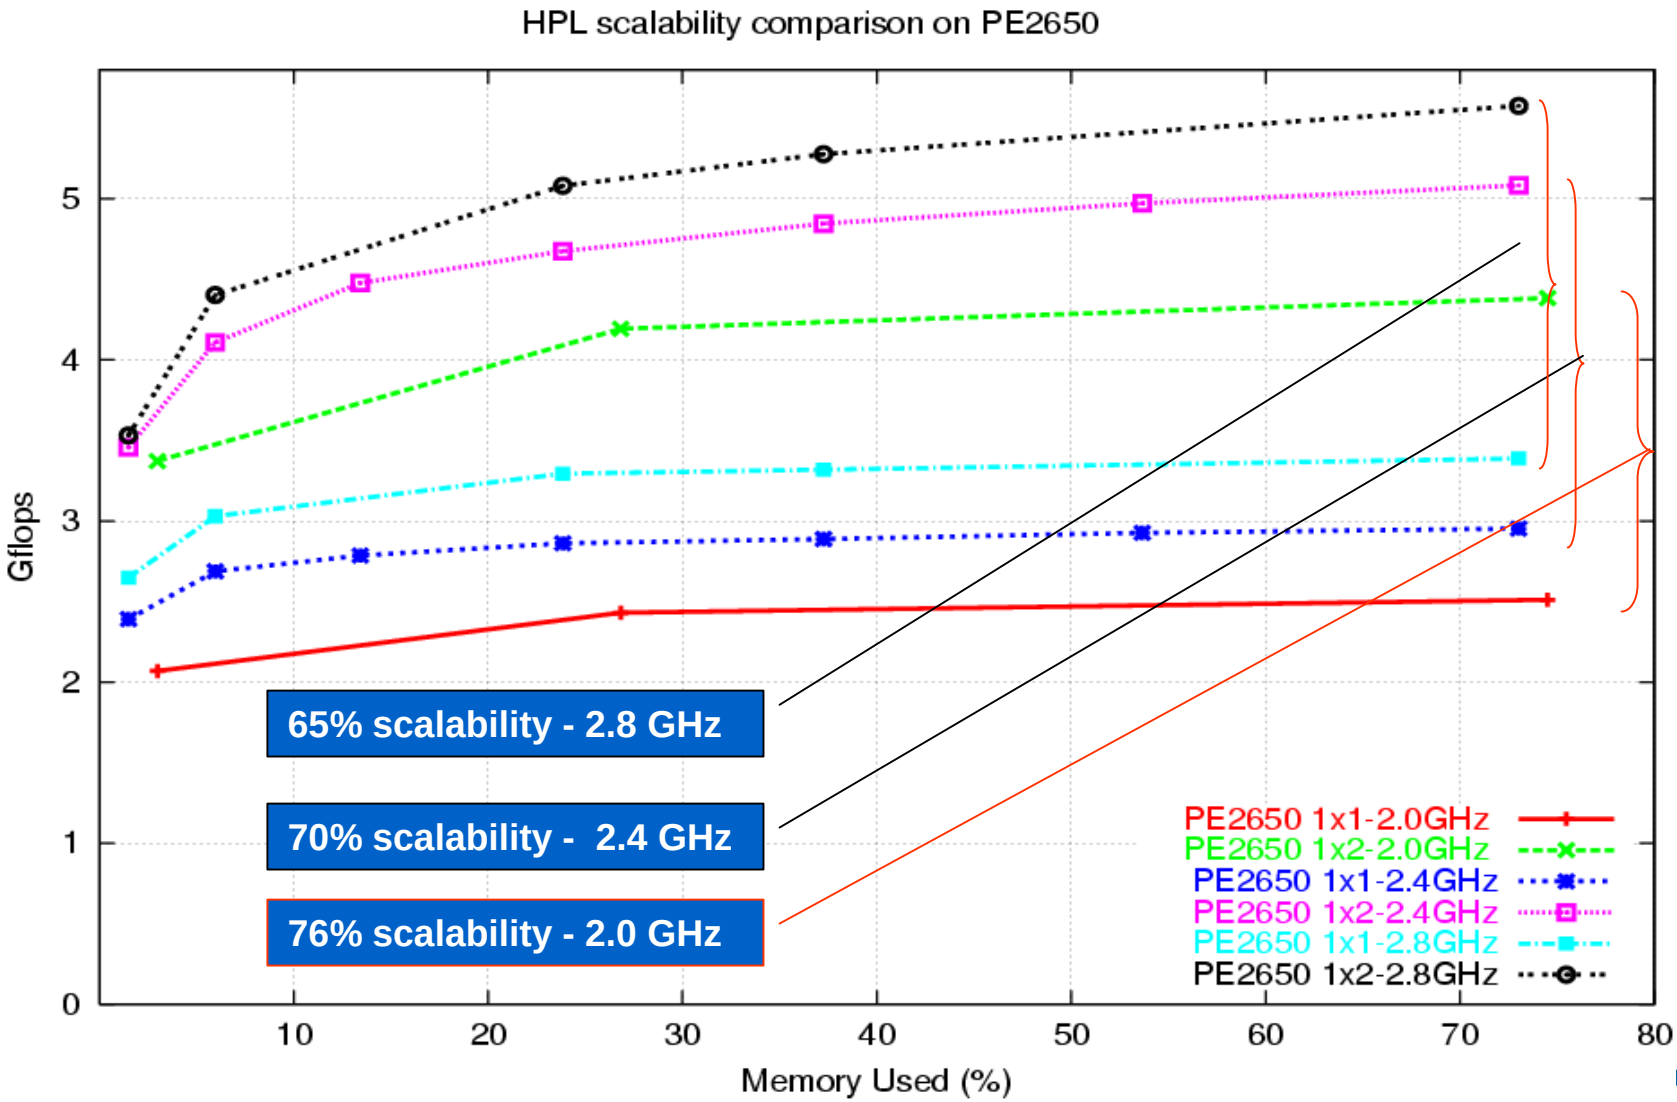

4-way vs. 2-way Interleaving using HINT

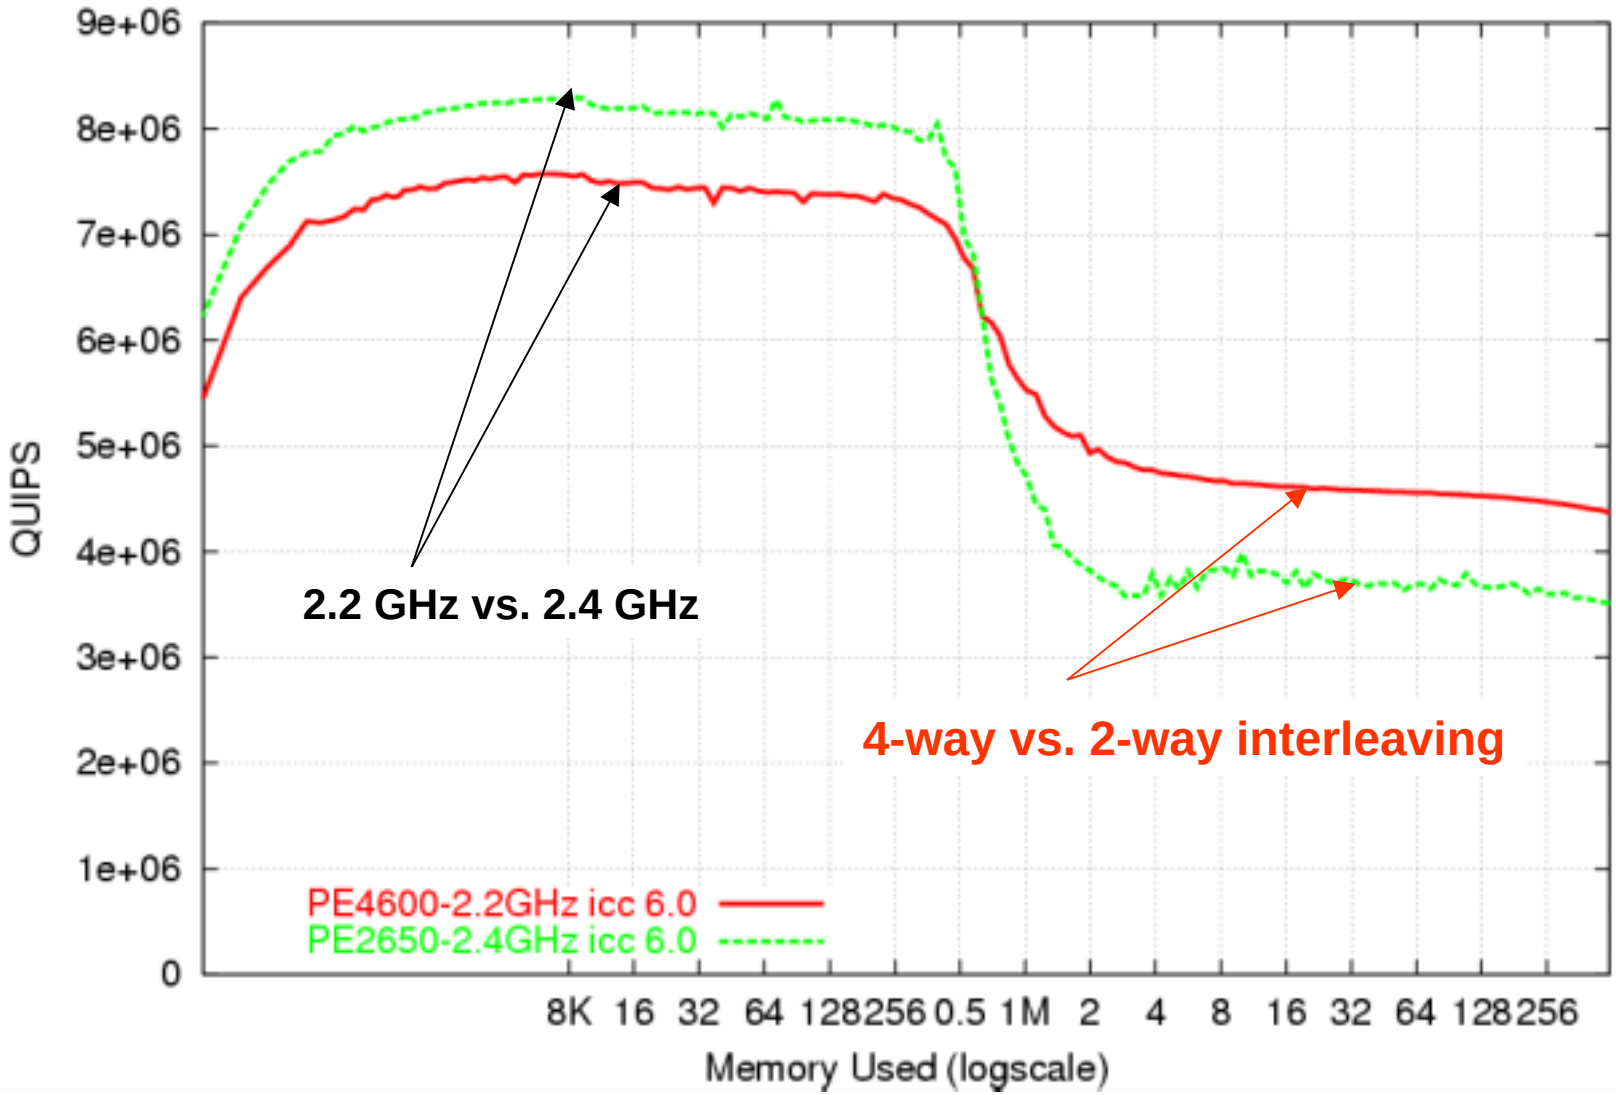

BLAST Performance Comparison

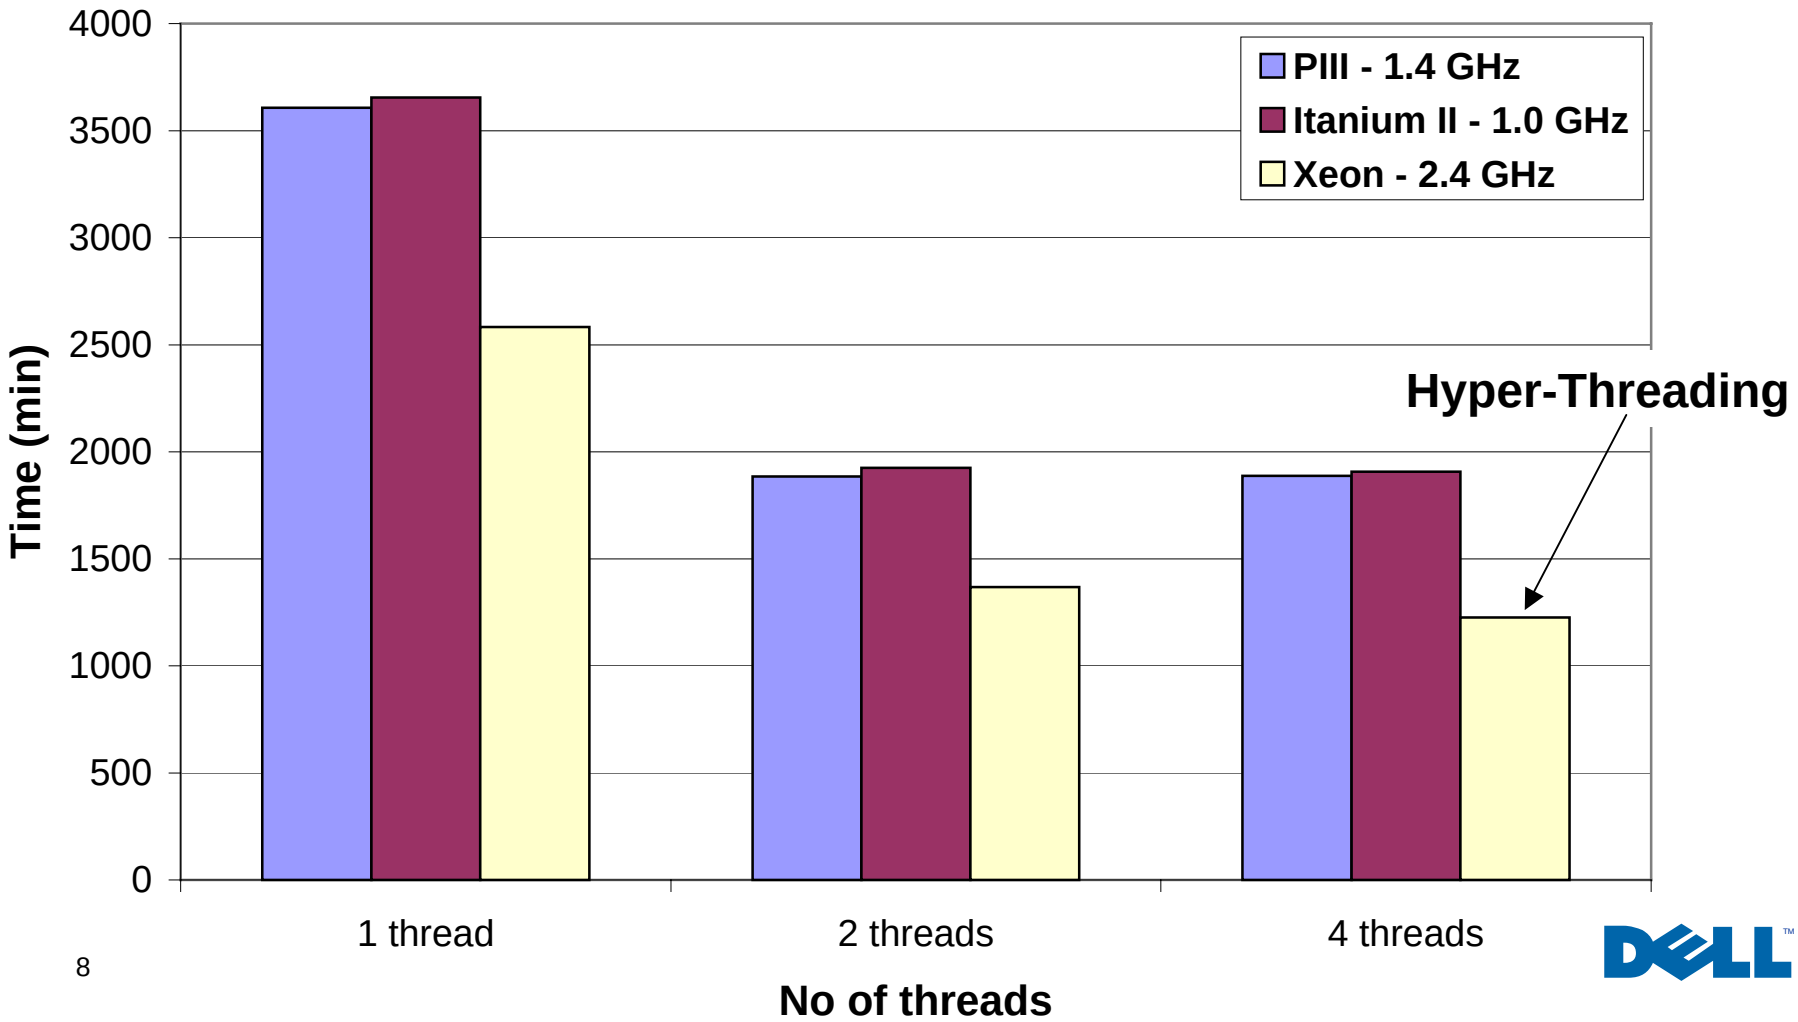

BLAS Comparison on Clusters

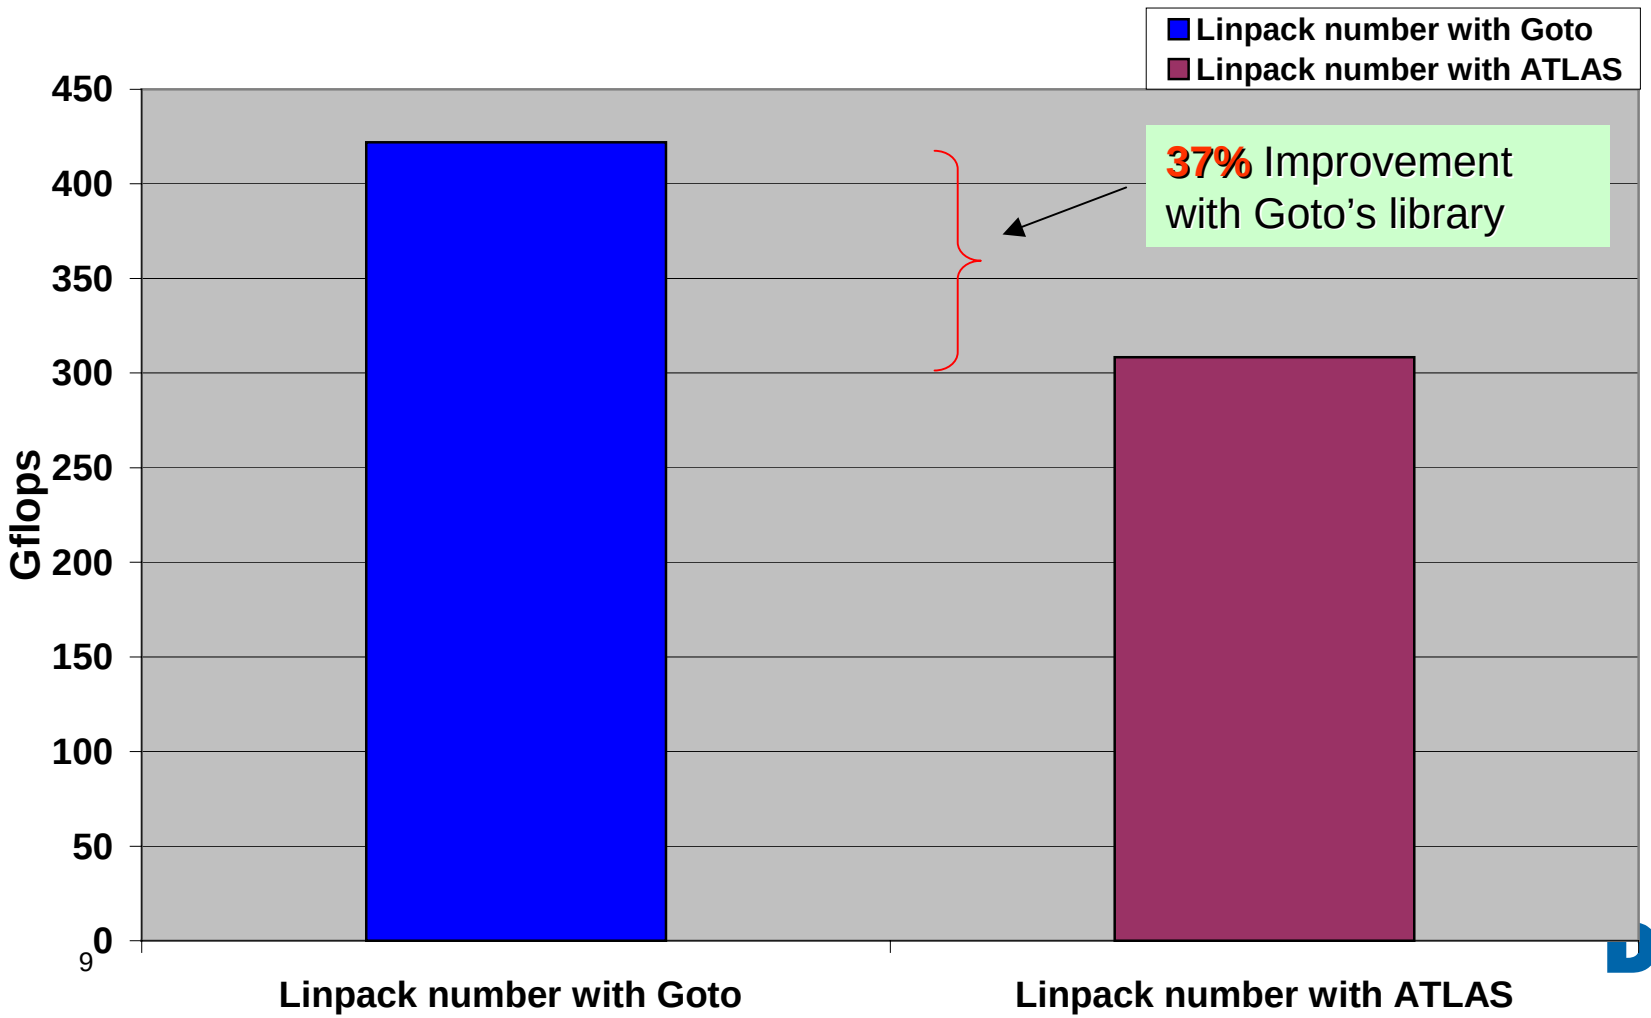

Aggregated Write Bandwidth

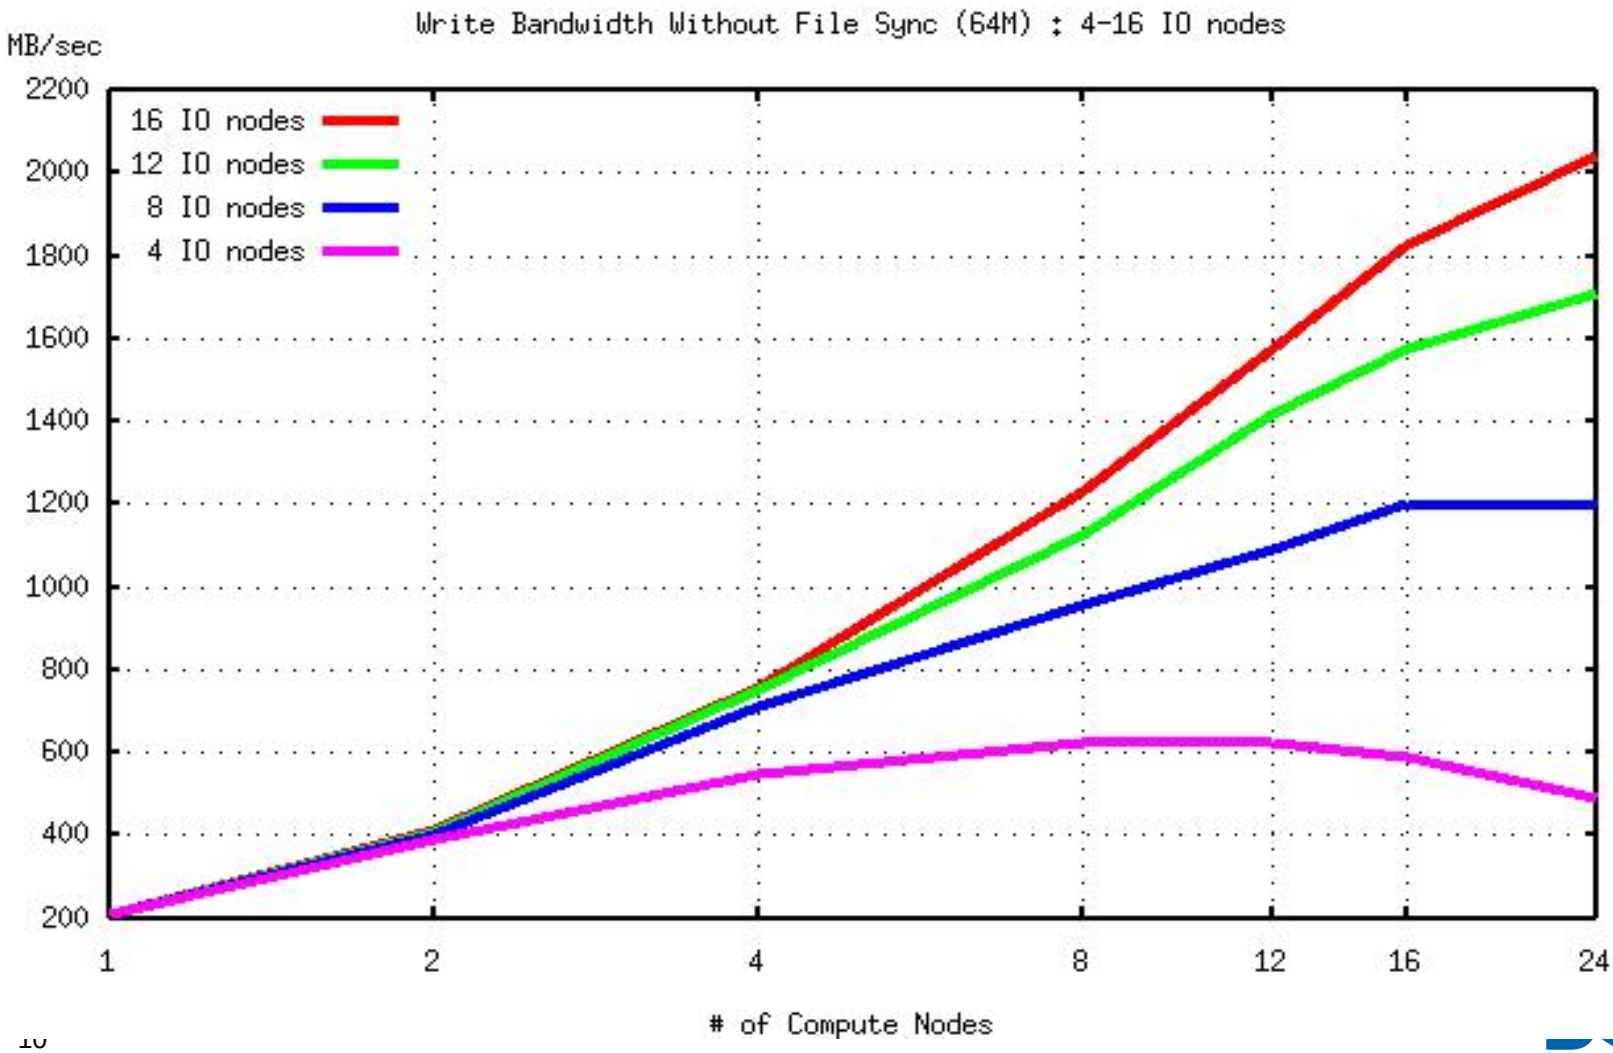

One Million Cell, Implicit, Black-Oil Model

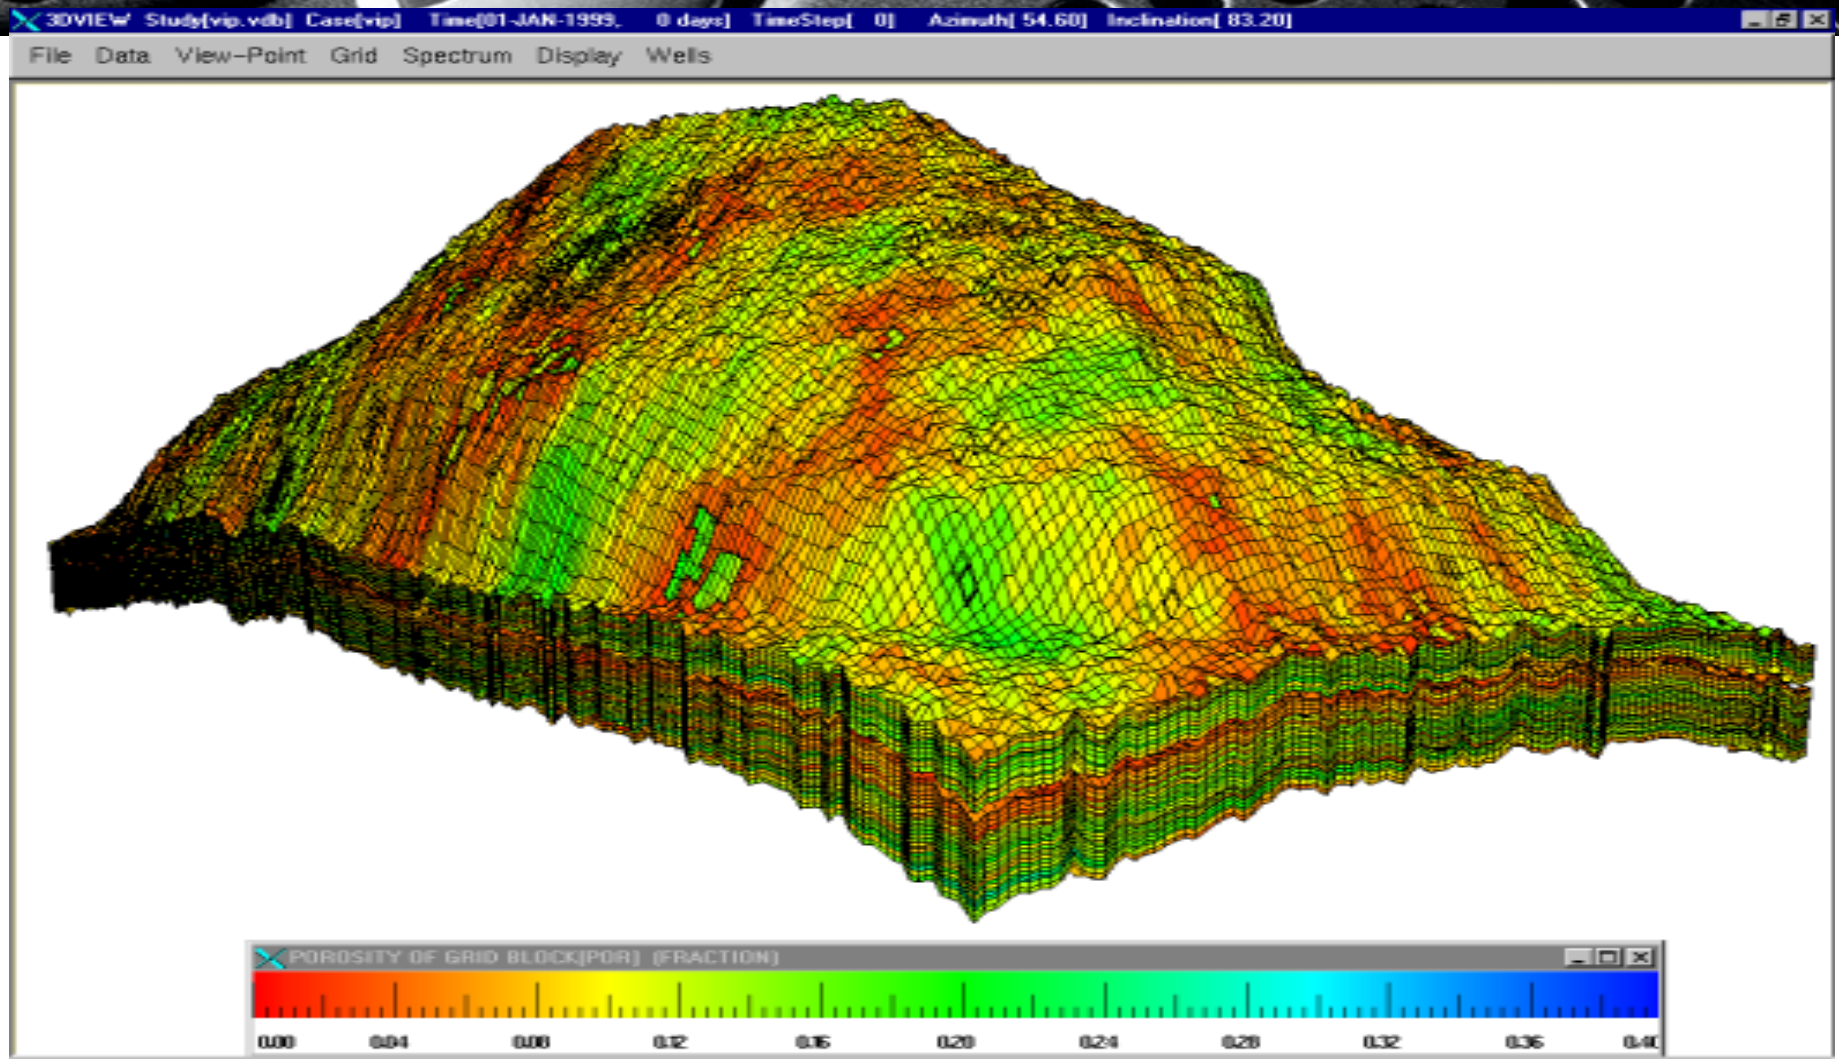

Source: Landmark Graphics

Price/Performance Comparison
UNIX vs. Xeon Clusters (W2K or LINUX)

1 Million Cell Model

CPU Type	Max CPU Time	Elapsed Time
8 Processor UNIX Machine A	320 seconds	355 seconds
8 Processor UNIX Machine B	220 seconds	221 seconds
8 2.4 GHz W2K Processors	92 seconds	100 seconds
16 2.4 GHz W2K Processors	56 seconds	63 seconds
8 2.2 GHz LINUX Processors	95 seconds	95 seconds
16 2.2 GHz LINUX Processors	67 seconds	67 seconds

Price Comparison

16 Processor/8 Gbyte Unix	US$300,000
16 Processor/8 Gbyte W2K Cluster	US$50,000

Source: Landmark Graphics

Sample of Dell HPCC Partners

- **Integration/Consultants**
  - Cray
  - MPI Software Technology, Inc
  - Cornell Theory Center
  - Scali
  - SCS
  - TurboWorx
- **Universities and National Lab**
  - Georgia Tech, College of Computing
  - Oak Ridge National Lab
  - Penn State University
  - University of Texas: Center of Petroleum & Geo-Systems Engineering
  - University of Houston Computer Science, High-performance Compilers
- **OS/Management Tools/ISV's**
  - CGG
  - Platform Computing
  - Fluent
  - Landmark Graphics
  - MSC.Software
  - Intel: Compilers
  - Microsoft
  - RedHat
- **Hardware Partners**
  - Intel
  - Myricom
  - Extreme Networks

Product Offerings

- Two classes of HPCC products: Standard and Custom
- Standard:
  - Low to medium size opportunities whose requirements can be generalized and packaged
  - To date, we have pre-tested/validated configurations of 8, 16, 32, 64 and 128 node configurations
  - Supports PIII and XEON technologies both
  - Fast and Gigabit Ethernet and Myrinet for intra-cluster communication
  - Fast Ethernet for management fabric
  - Software stack for building a generic HPC stack
  - Professional services from pre-sales to post-sales
- Custom:
  - Case-by-case, larger or strategic opportunities that have unique customer requirements and have to be handled individually.

Dell Centers for Research Excellence Awards

- Award created by Michael Dell to recognize innovative uses of High Performance Compute Clusters
  - Innovation in HPCC applications or solutions: organizations that develop technical enhancements that further the standardization and simplify the use of cluster computing for data intensive applications.
  - Size and scope of the cluster: organizations that have HPCC deployments that achieve new levels of performance and capabilities.
  - Applications and types of research: organizations that use a HPCC cluster to perform groundbreaking commercial and government research or research for the betterment of society.

Technical Computing Market

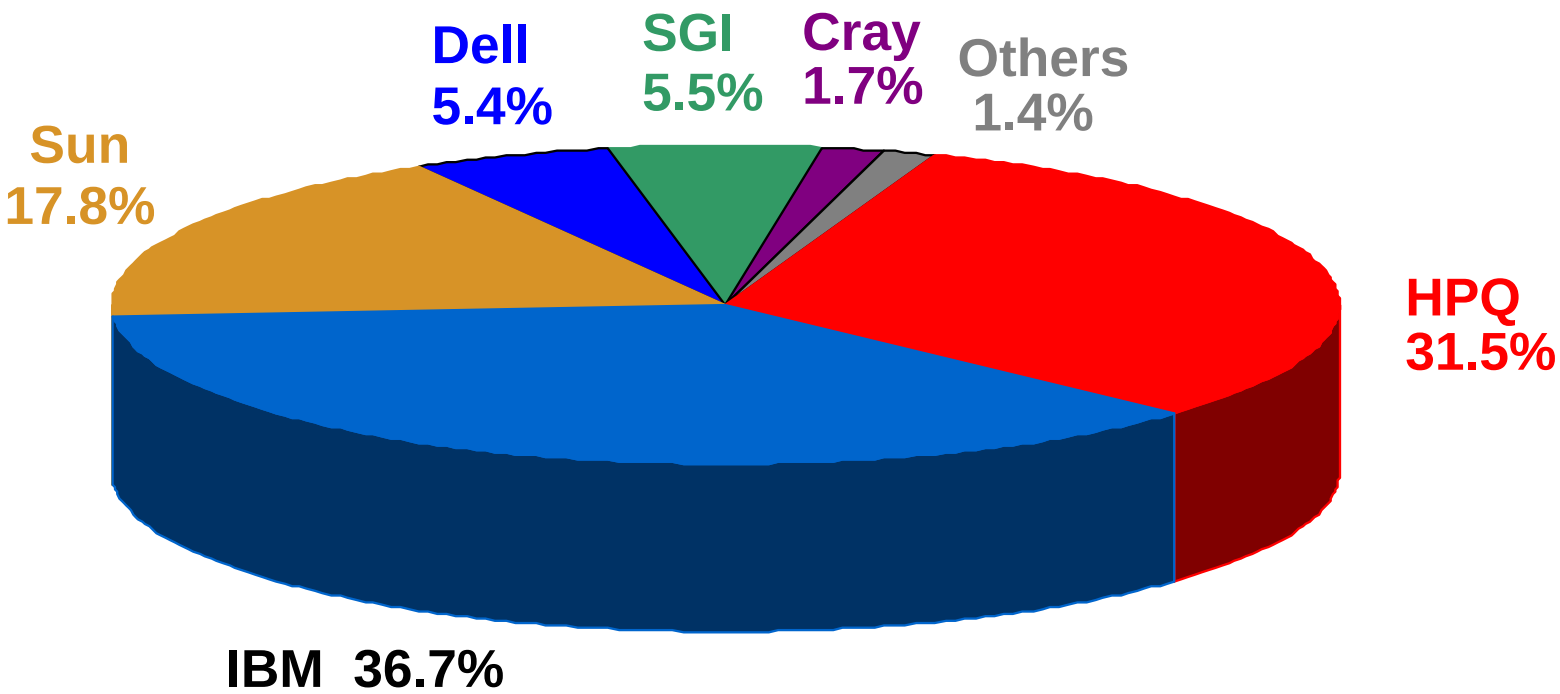

CY'01 Dell was part of the "Others" group: We have a Great deal of work left to do!

Source: IDC High Performance Technical Computer QView, Q4'02

Thank you for your time!

Questions?