The Supercomputer Company

Taking the Lead in HPC


Cray X1


Cray XD1

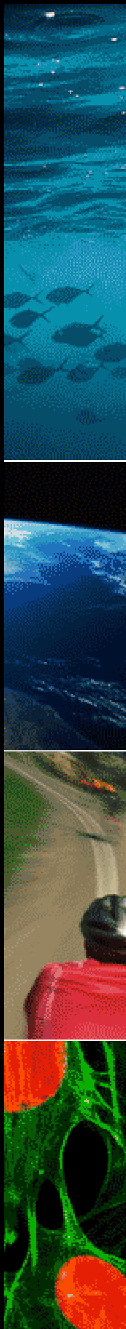

Cray Update
SuperComputing 2004

Safe Harbor Statement

The statements set forth in this presentation include forward-looking statements that involve risk and uncertainties. The Company wished to caution that a number of factors could cause actual results to differ materially from those in the forward-looking statements. These and other factors which could cause actual results to differ materially from those in the forward-looking statements are discussed in the Company's filings with the Securities and Exchange Commission.

Agenda

- **Introduction & Overview – Jim Rottsolk**
- Sales Strategy – Peter Ungaro
- Purpose Built Systems - Ly Pham
- Future Directions – Burton Smith
- Closing Comments – Jim Rottsolk

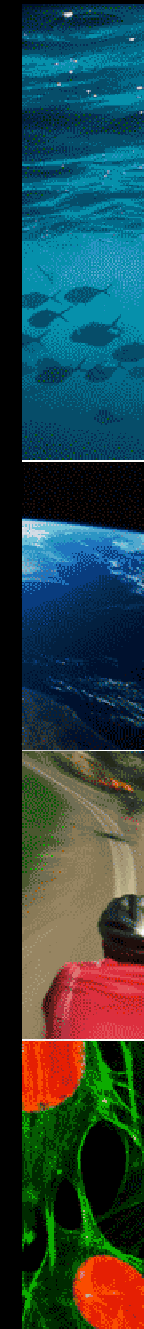

A Long and Proud History

NASDAQ:	CRAY
Headquarters:	Seattle, WA
Marketplace:	High Performance Computing
Presence:	Systems in over 30 countries
Employees:	Over 800

Building Systems Engineered for Performance

Seymour Cray
Founded Cray Research 1972
The Father of Supercomputing

Cray-1 System (1976)
First Supercomputer

Cray-X-MP System (1982)
First Multiprocessor Supercomputer

Cray Y-MP System (1988)
First GFLOP Computer

Cray T3E System (1996)
World's Most Successful MPP
First TFLOP Computer

Cray X1
8 of top 10 Fastest
Computer in the World

2002 to 2003 – Successful Introduction of X1

Market Share growth exceeded all expectations

- 2001 Market: $800M

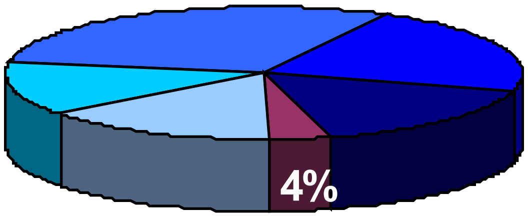

- 2002 Market: about $1B

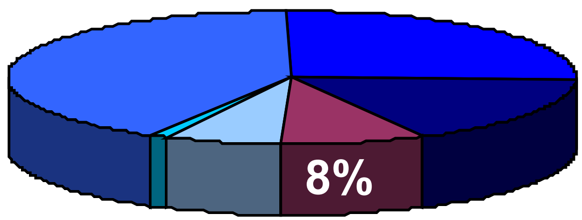

- 2003 Market: $770M

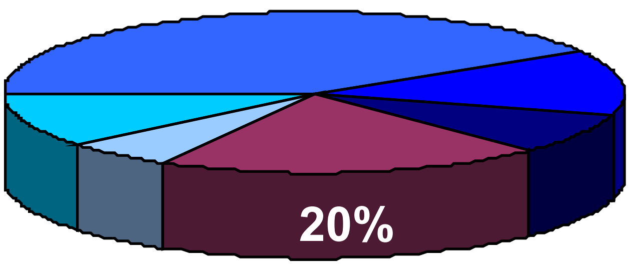

Source: IDC, 2003 Capability Market Census

2004: Transition & Restructuring

- One product company to complete HPC portfolio

Cray XD1

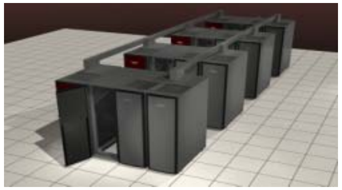

Cray XT3

Cray X1E

- Disappointing near-term financial results
  - **Drop-off in X1 business - particularly government**
  - **Push-out of new product introductions – X1E and XT3**
- Third quarter restructuring
- 2004 full-year revenue outlook: $155MM - $165MM
- Expect return to profitability in early 2005

2005 Outlook

- Growth Year as HPC Leadership Company
  - **Three highly anticipated new products in full-swing**
  - **Significant top and bottom-line growth opportunity**
  - **Positive EPS affect via leveraged business model & recent restructuring**
  - **Continued product innovation focused on high performance computing**

Cray X1E

Upgrade to Cray X1
Double performance & density
First shipment end of 2004

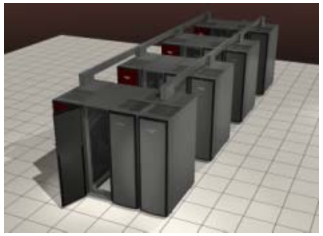

Cray XT3

Largest x86 system in the world
Based on contract with Sandia
First shipment Q4 2004

Cray XD1

Scales to over 512 processors
Supercomputer from under $100K
General Availability Oct. 2004

Market Opportunity

Taking Success Formula to the Broader HPC Market.

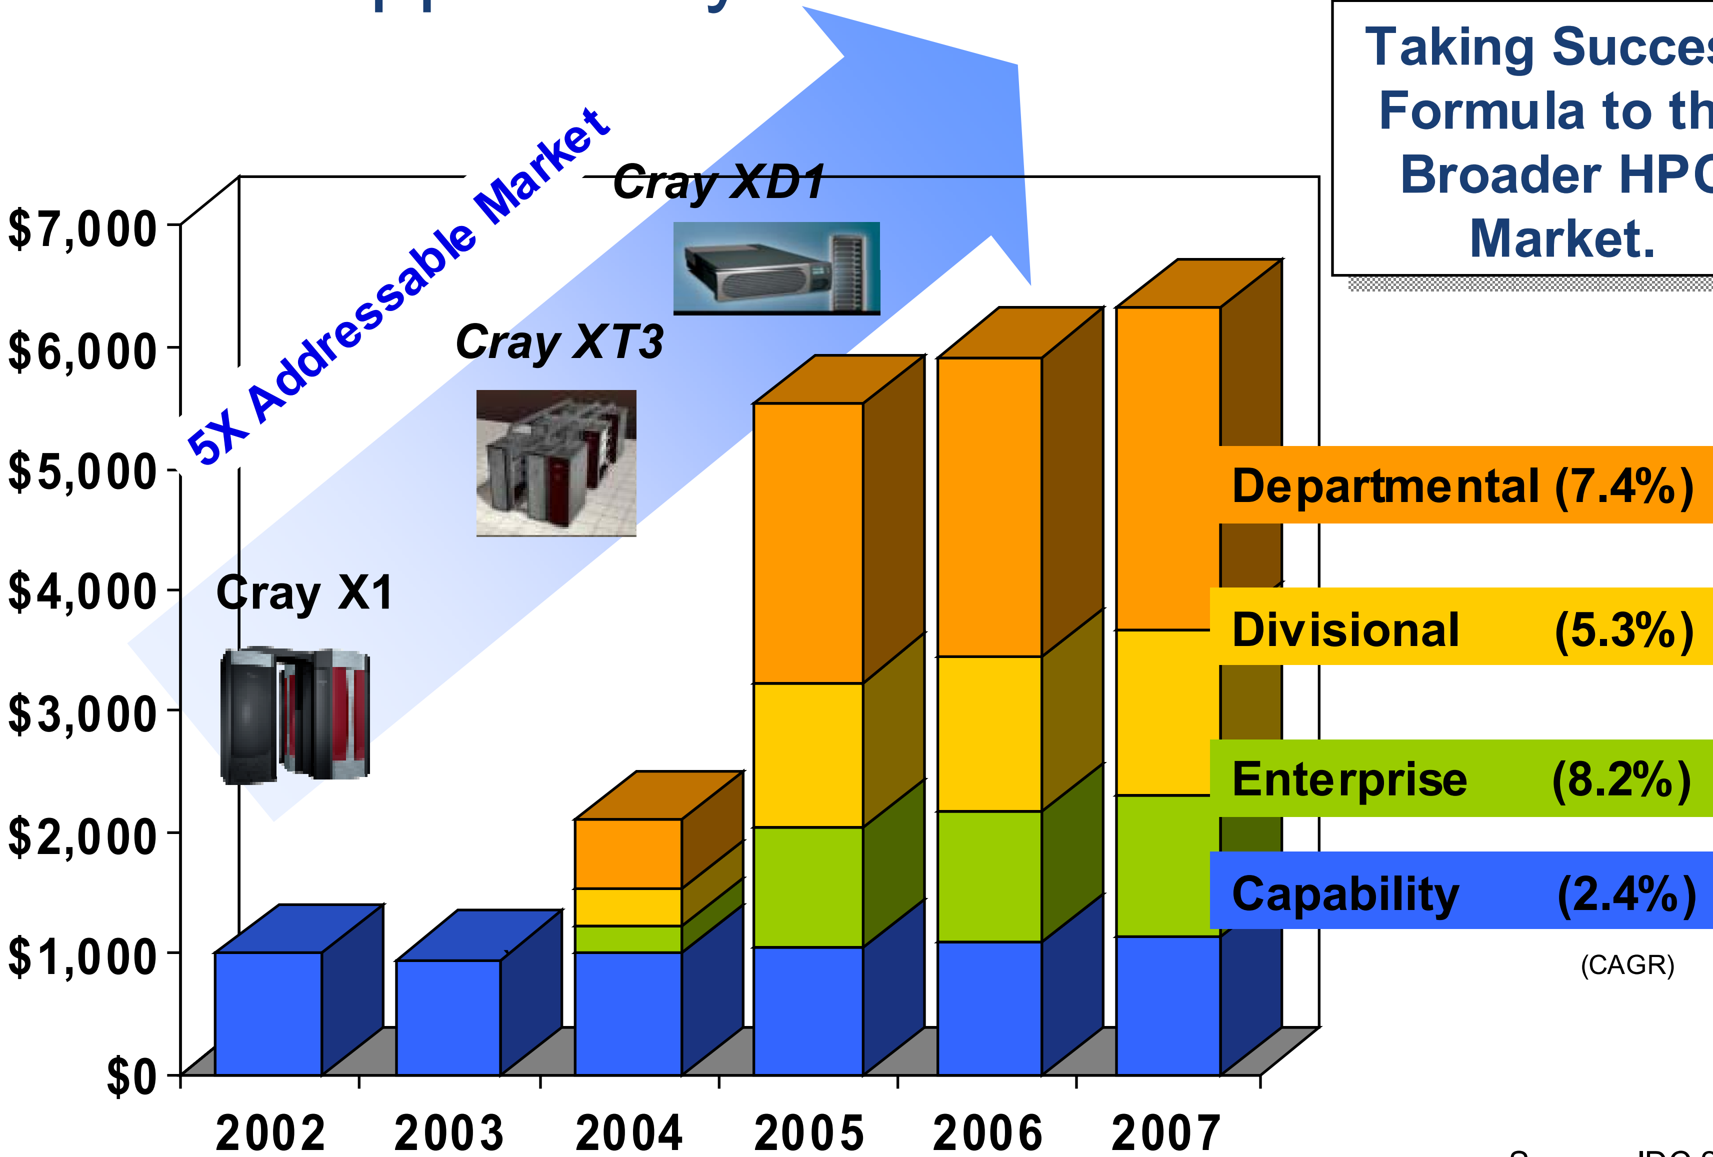

Source: IDC 2003

Agenda

- Introduction & Overview – Jim Rottsolk
- **Sales Strategy – Peter Ungaro**
- Purpose Built Systems - Ly Pham
- Future Directions – Burton Smith
- Closing Comments – Jim Rottsolk

The Supercomputer Company

Sales & Marketing 2004 – 2005

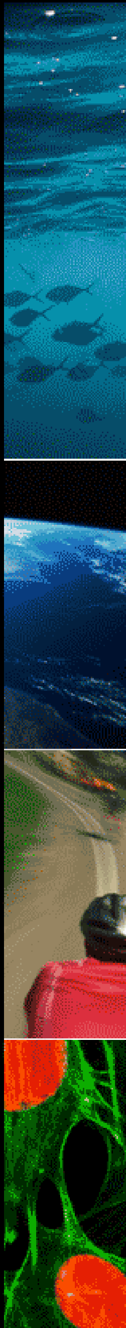

Peter Ungaro
Senior Vice President
Sales, Marketing and Services

Diverging Market Approaches

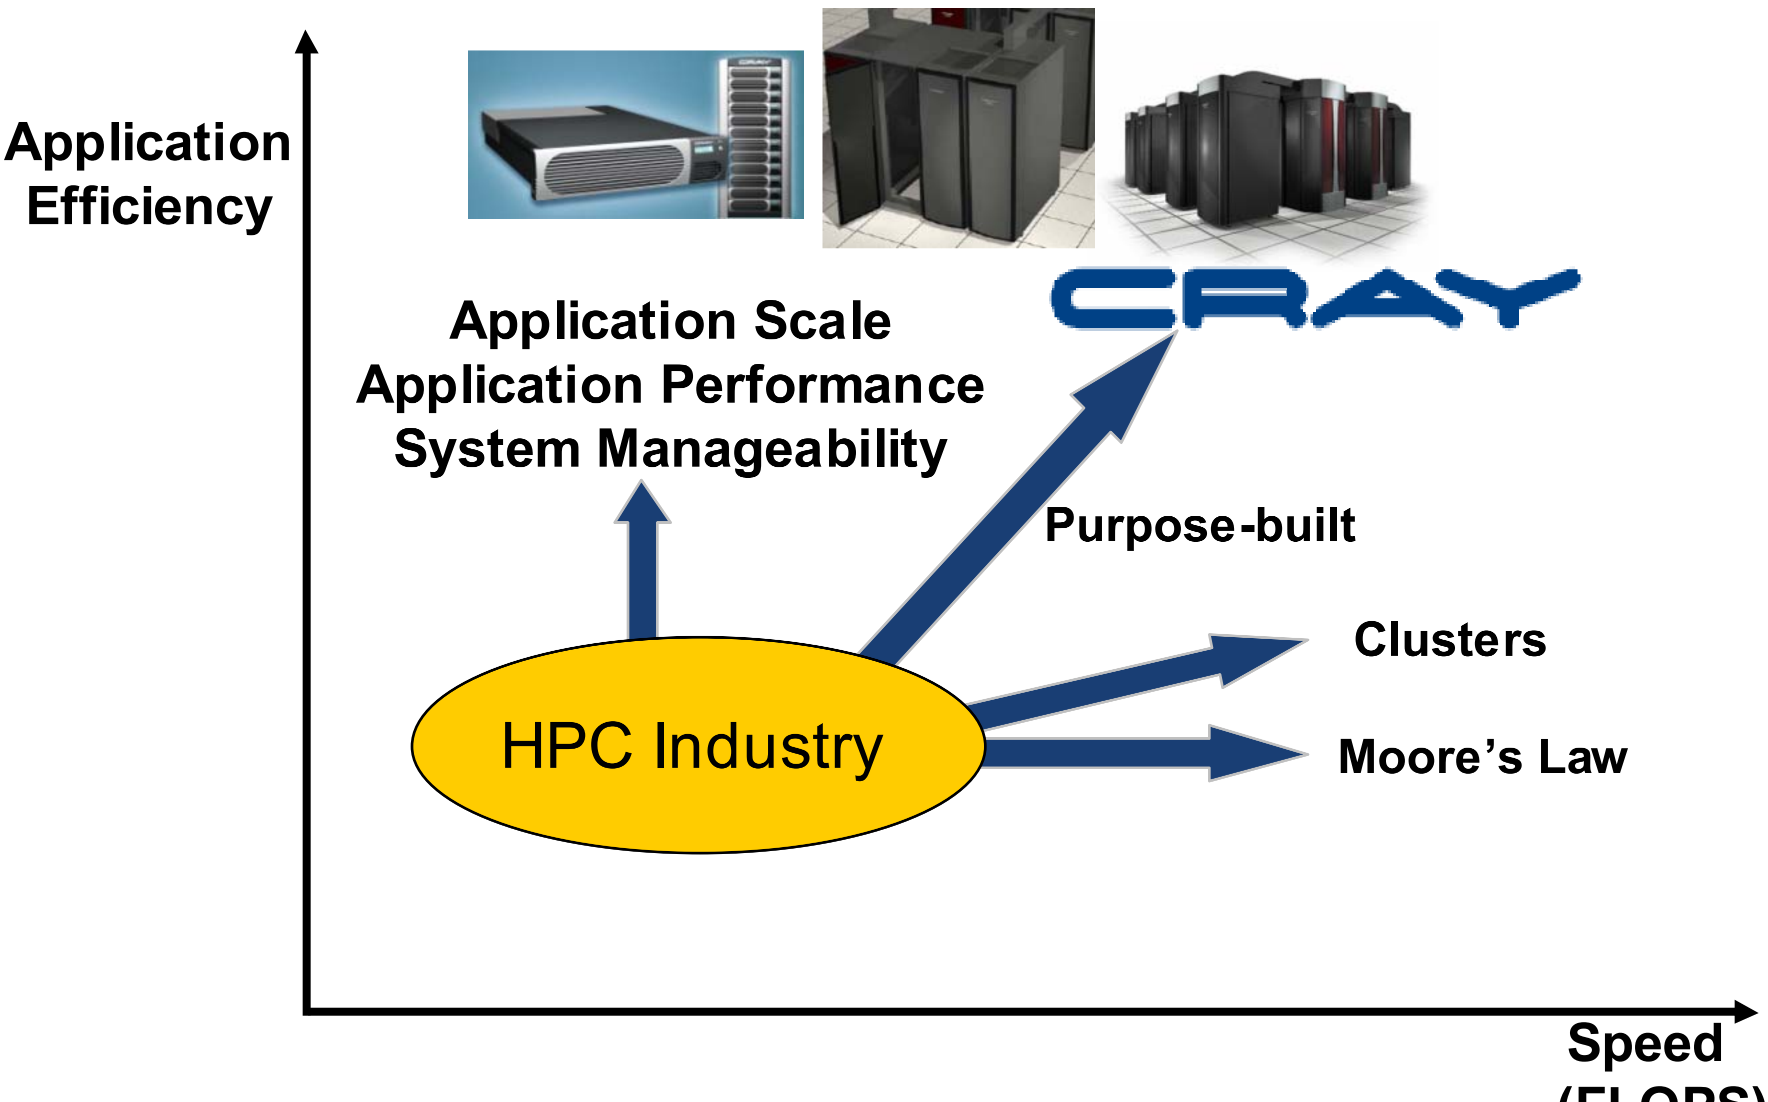

Entering 2004

Cray X1

- New wins – KMA - Oak Ridge
- Upgrades from Cray X1 to X1E in 2005.
- Fastest system on HPC Challenge

Bringing purpose-built HPC solutions to High End HPC

Comprehensive Cray Portfolio

Cray X1

- New wins – KMA - Oak Ridge
- Upgrades from Cray X1 to X1E in 2005.
- Fastest system on HPC Challenge

Cray XT3

- Shipped first installment of Red Storm to Sandia
- Sales to Oak Ridge National Labs, Pittsburgh Supercomputing Center plus other, undisclosed customers.
- Winner of 5 HPCWire Readers Choice awards

Cray XD1

- General Availability Oct 4
- Customers in every major geography
- Lowest latency interconnect in the industry.
- Demonstrated performance advantage in CAE, computational chemistry, weather & climate modeling, reconfigurable computing

Customers across the board turning to Cray for Purpose-built HPC solutions

Comprehensive Cray Portfolio

Cray X1

Vector Processors for Uncompromised Sustained Performance

Cray XT3

MPP Compute System for Large-Scale Sustained Performance

Cray XD1

Entry/Mid Range System Optimized for Sustained Performance

Addressing the high bandwidth needs of the entire HPC market

Comprehensive Cray Portfolio

Cray X1

- 1 to 50+ TFLOPS
- $3 M+
- Vectorized apps
- Cray MSP with UNICOS/mp

Cray XT3

- 1 to 50+ TFLOPS
- 256 – 10,000+ processors
- $1 M+
- AMD Opteron with UNICOS/lc

Cray XD1

- 58 GFLOPS – 2.5+ TFLOPS
- 12 – 512+ processors
- $50 K+
- AMD Opteron with Linux

Complementary solutions for broad HPC market requirements

Leadership Class Computing

OAK RIDGE NATIONAL LABORATORY
CCS The Center for Computational Sciences
DOE High Performance Computing Research Center

- **Cray-ORNL Selected by DOE for National Leadership Computing Facility (NLCF)**
- **Goal: Build the most powerful supercomputer in the world**
- **250-teraflop capability by 2007**
  - **50-100 TF sustained performance on challenging scientific applications**
  - **Roadmap:**
    - **512p MSP X1**
    - **20T X1E**
    - **20T XT3**
    - **100T Rainier**
- **Focused on capability computing**
  - **Available across government, academia, and industry**
  - **Including biology, climate, fusion, materials, nanotech, chemistry**
  - **Open scientific research**

Shipped first Red Storm System Installment

Sandia National Laboratories

- 41.5TF peak performance
- 140 cabinets
  - 11,648 AMD Opteron™ processors
  - 10 TB DDR memory
- 240 TB of disk storage
- Approximately 3,000 ft²

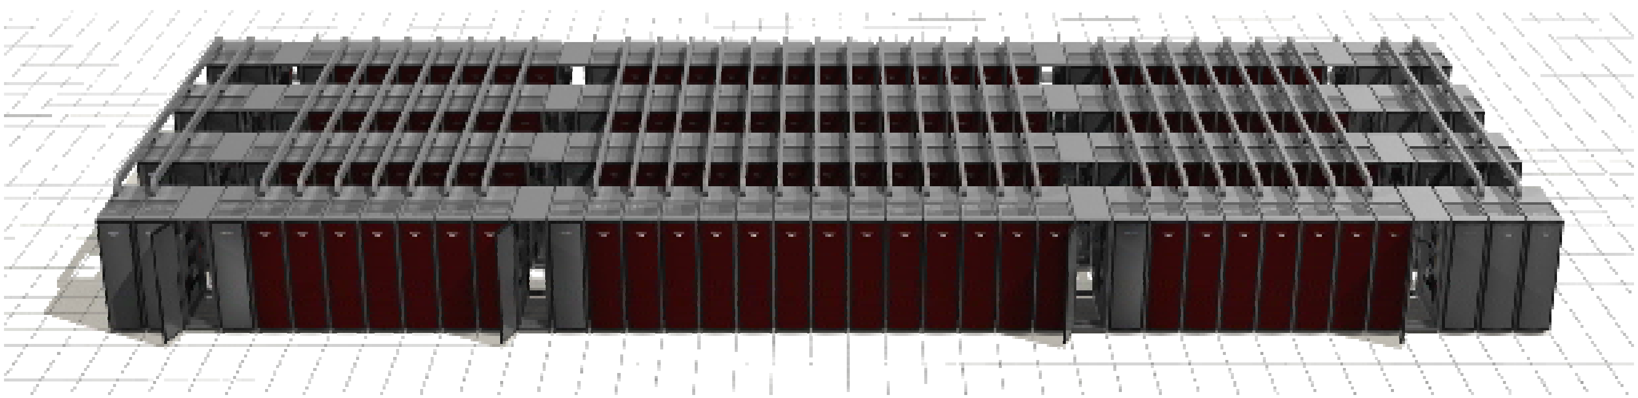

"We expect to get substantially more real work done, at a lower overall cost, on a highly balanced system like Red Storm than on a large-scale cluster." Bill Camp, Sandia Director of Computers, Computation, Information and Mathematics

Pittsburgh Supercomputer Center

PITTSBURGH SUPERCOMPUTING CENTER

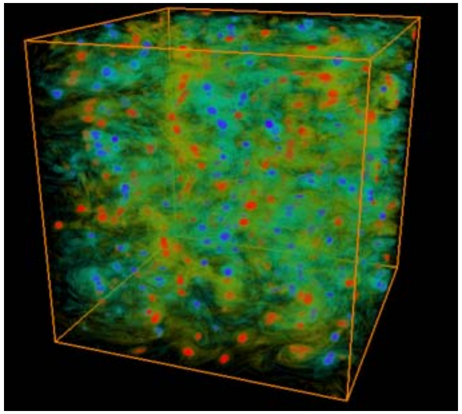

Environmental sciences: Global climate modeling

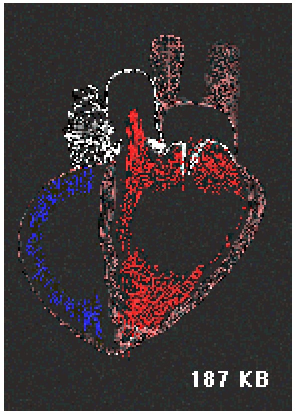

CAE: blood flow modeling

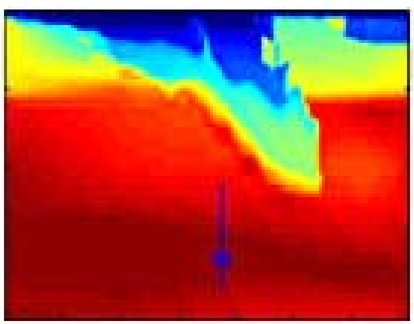

Environmental sciences: Earthquake ground vibration

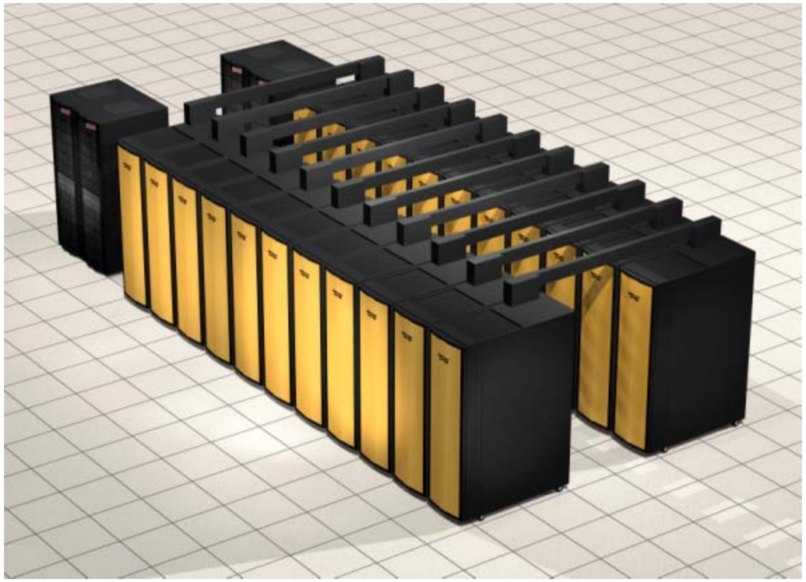

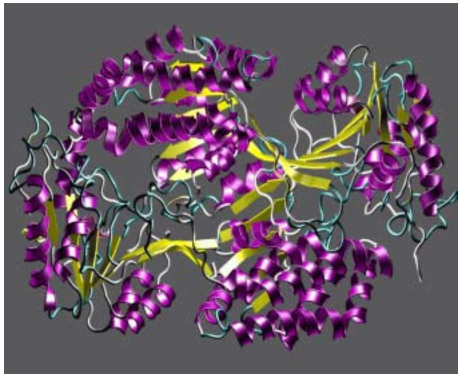

Life sciences: protein simulation

10 TFLOPS Cray XT3 System
2,000 AMD Opteron Processors

US Forest Service selects Cray XD1

Fire Sciences Lab
PO Box 8089 - 5775 W Hwy 10 - Missoula MT 59807

- Cray XD1 supercomputer will help US Forest Service predict behavior of forest fires
- Application: WRF-Chem
- Requirement:
  - 48 hour forecast across the entire US at 12 km resolution with two hours of run time.
- Selection criteria
  - Price/performance of test case
  - Cost of future expansion
  - Provision for and cost of scheduled maintenance
- Benchmark result: XD1 will complete US Forest Service WRF-chem forecast in about 1 hour 6 minutes on 24 XD1 processors

Winning performance puts Cray XD1 in US Forest Service Fire Sciences Lab

Researchers exploiting FPGAs

- Ohio Supercomputing Center
  - Establishing Reconfigurable Computing center of expertise
  - Installed Cray XD1 supercomputer in August
  - Technical Symposium on Reconfigurable Computing with FPGAs, October 4-5, 2004, Springfield, Ohio
- PNNL and other leading research centers selected Cray XD1 for FPGA
- "easiest to use FPGA system we've seen"

Advancing reconfigurable computing through open source development

Building the Infrastructure: People, process, partnerships

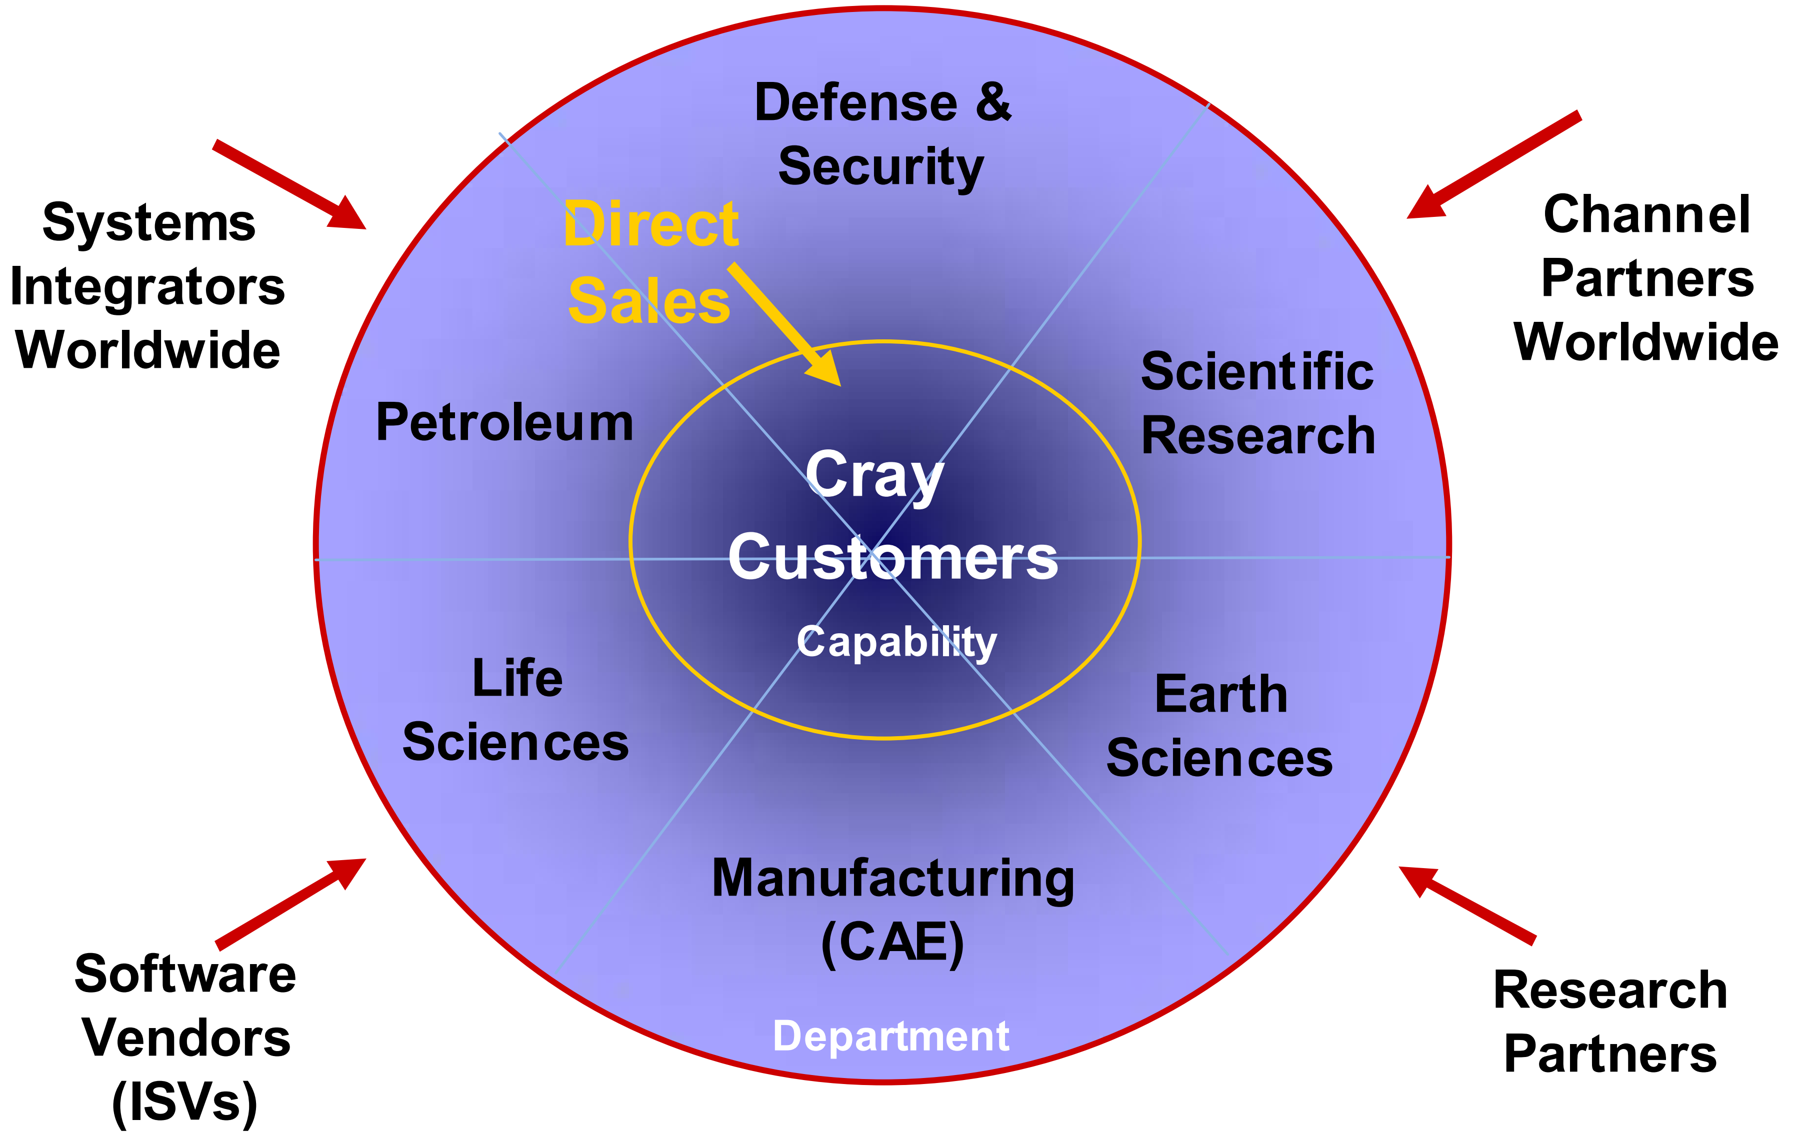

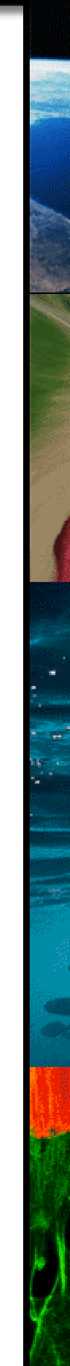

Agenda

- Introduction & Overview – Jim Rottsolk
- Sales Strategy – Peter Ungaro
- **Purpose Built Systems - Ly Pham**
- Future Directions - Burton Smith
- Closing Comments – Jim Rottsolk

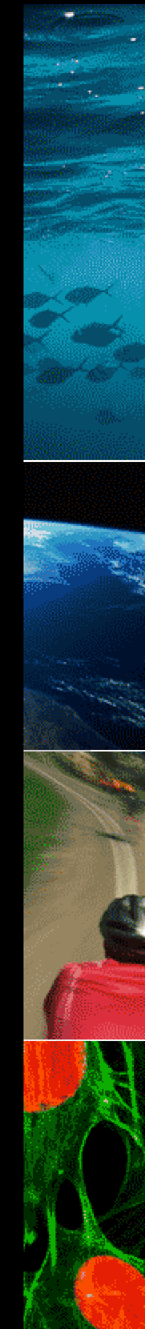

CRAY The Supercomputer Company

Supercomputers: Systems purpose-built for the demands of HPC

Ly Pham
Senior Vice President
Operations

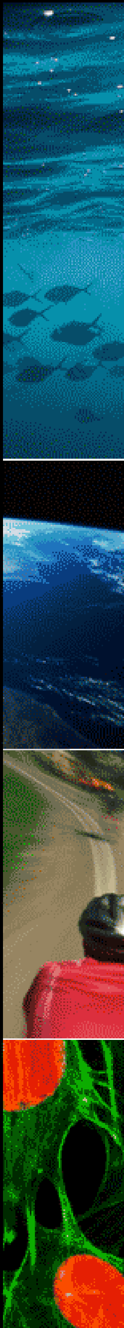

A Brief Bio

Ly Pham

turbolinux®

VTEL Products Corporation

PEPPERDINE UNIVERSITY
Graziadio School of Business and Management

WANG

BOSTON UNIVERSITY

Our Customers' Problems

Manufacturing

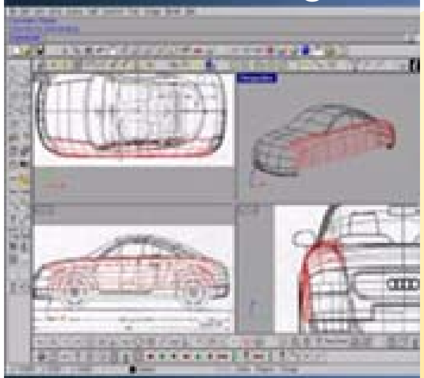

Modeling new designs

Life Sciences

Drug discovery

Government

Predicting weather

HPC Computing Challenge

Scale

Performance

Manageability

Technology Limitations

Interconnect Bottleneck

Can't manage as a single system

Manageability

Oil & Gas

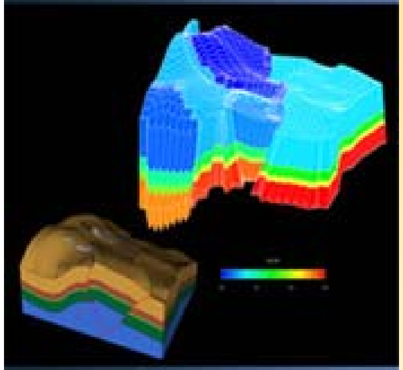

Exploring the Earth

Scientific Research

Exploring the Universe

Purpose-Built for HPC

Vector Processing

High Bandwidth Computing

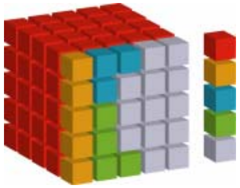

Scalable Software

Reliability, Availability, Manageability

HPC Expertise that Permeates the Company

Vector Processing

Cray X1 Vector Processor

- Powerful vector processors
  - **Over 12.8 GFLOPS (2-5x commodity processors)**
  - **205 GB/s peak bandwidth to local memory (30x commodity systems)**
- Highly-scalable network interconnect
  - **Low latency, high bandwidth (34.1 GB/sec; 5-130x faster than commodity interconnects)**

Vector Processors for Uncompromised Sustained Performance

Removing the Communications Bottleneck

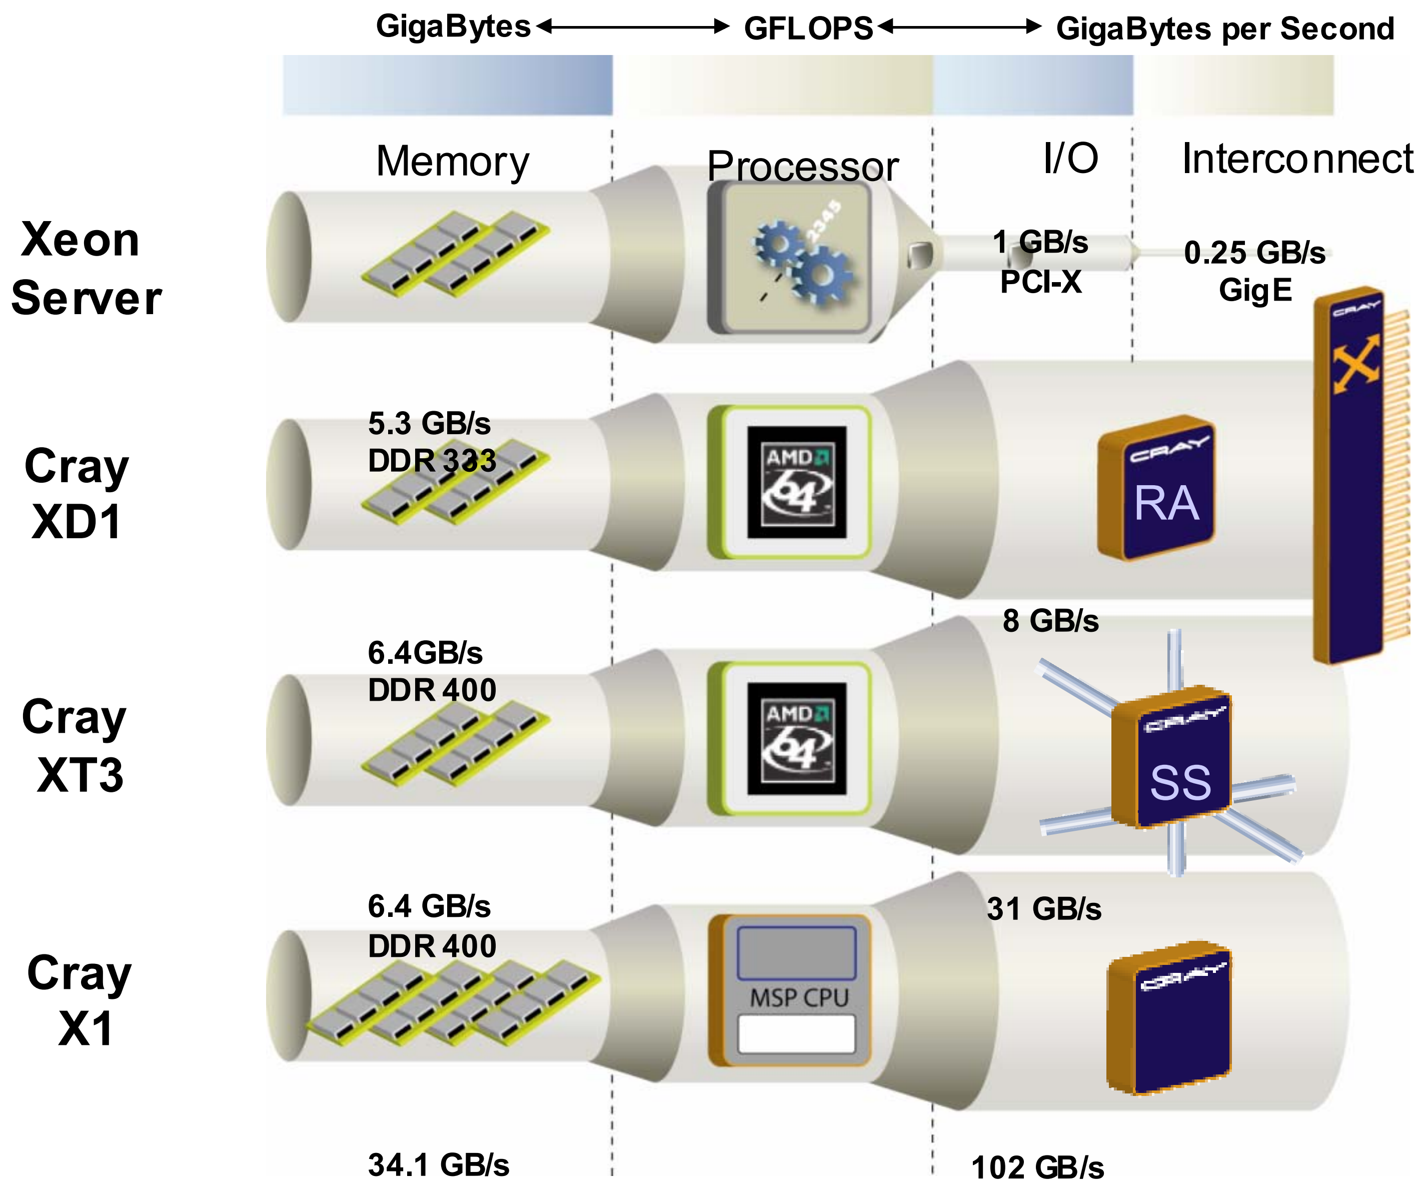

Scalable Software

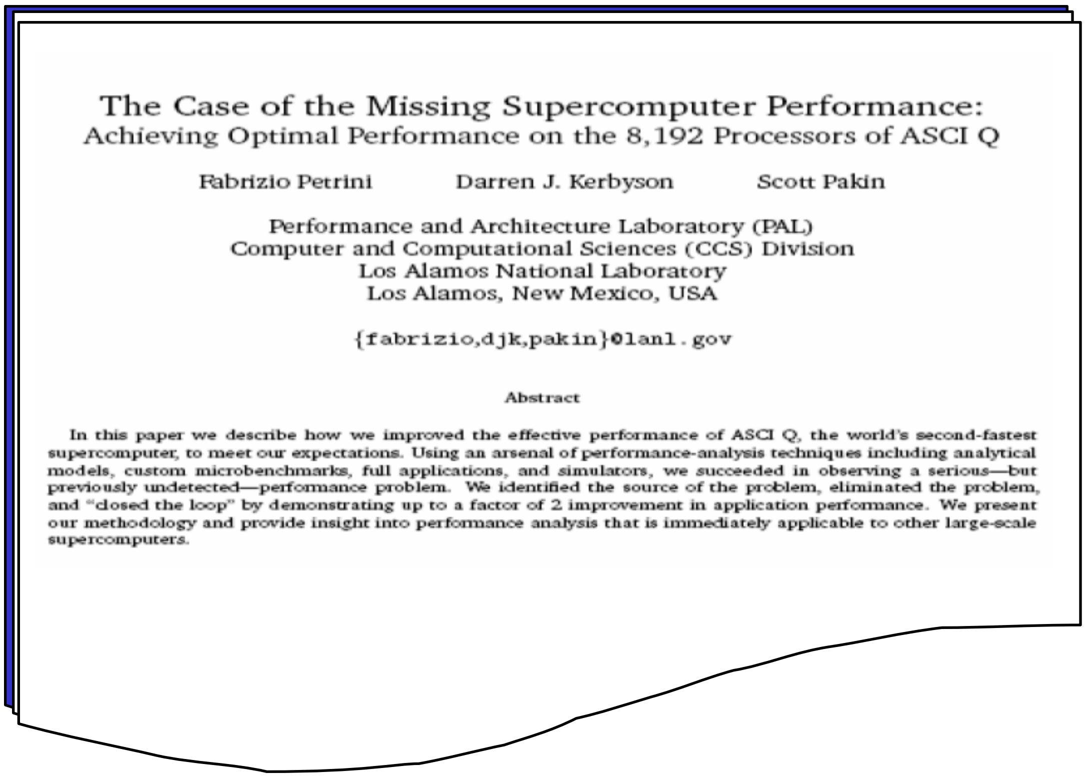

The Case of the Missing Supercomputer Performance:
Achieving Optimal Performance on the 8,192 Processors of ASCI Q

Fabrizio Petrini   Darren J. Kerbyson   Scott Pakin

Performance and Architecture Laboratory (PAL)
Computer and Computational Sciences (CCS) Division
Los Alamos National Laboratory
Los Alamos, New Mexico, USA

{fabrizio,djk,pakin}@lanl.gov

Abstract

In this paper we describe how we improved the effective performance of ASCI Q, the world's second-fastest supercomputer, to meet our expectations. Using an arsenal of performance-analysis techniques including analytical models, custom microbenchmarks, full applications, and simulators, we succeeded in observing a serious—but previously undetected—performance problem. We identified the source of the problem, eliminated the problem, and "closed the loop" by demonstrating up to a factor of 2 improvement in application performance. We present our methodology and provide insight into performance analysis that is immediately applicable to other large-scale supercomputers.

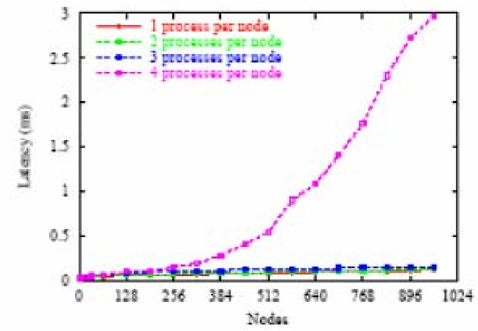

Doubling Application Performance

Reliability, Availability, Manageability

IDC Cluster User Study – 2004

"The biggest challenge to implementing clusters is systems management (43%)"

"Users are primarily concerned about the usability, manageability and long term operation of installed cluster systems"

"One cluster user described the management problem this way: All the components work, they just don't always work together" (IDC 2002)

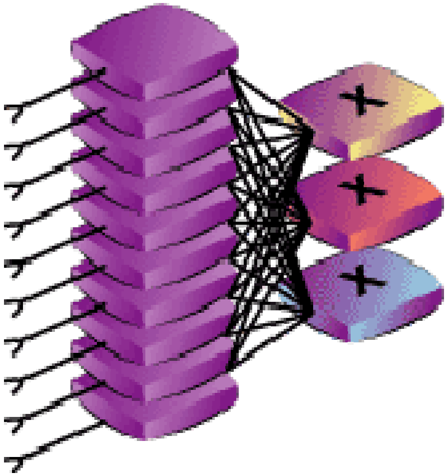

It is crucial to address the management issues

Cray's Computing Vision

Scalable High-Bandwidth Computing

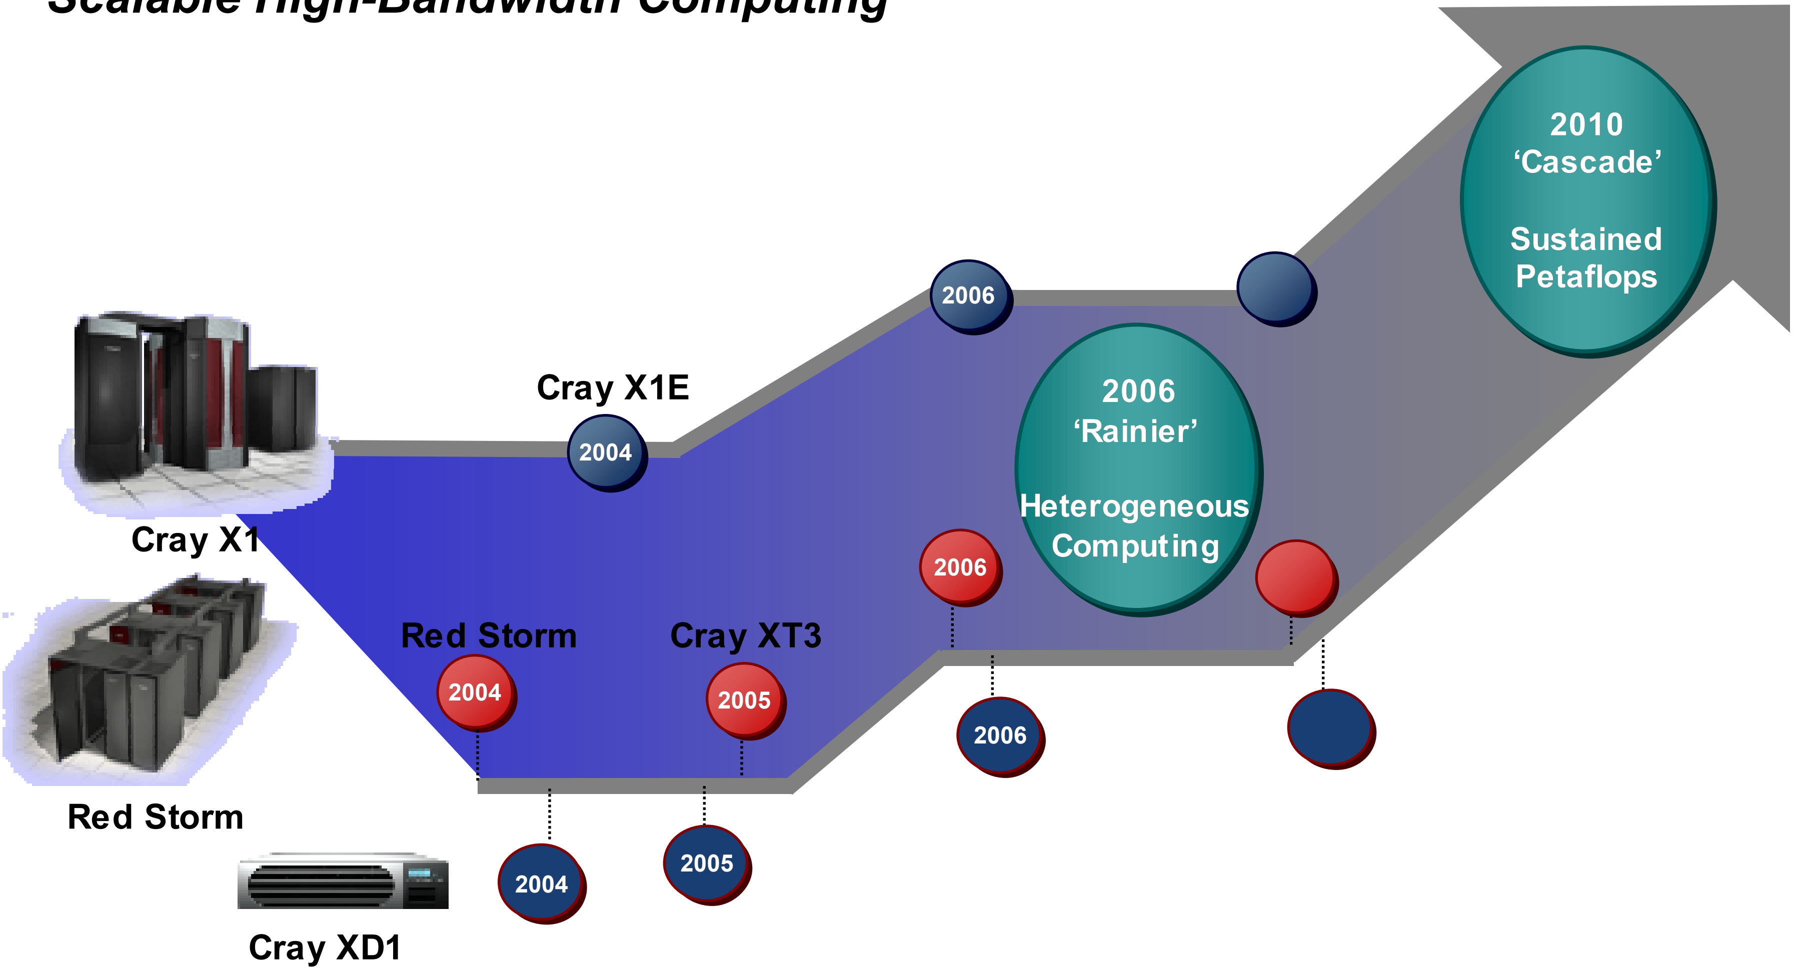

Building the Products

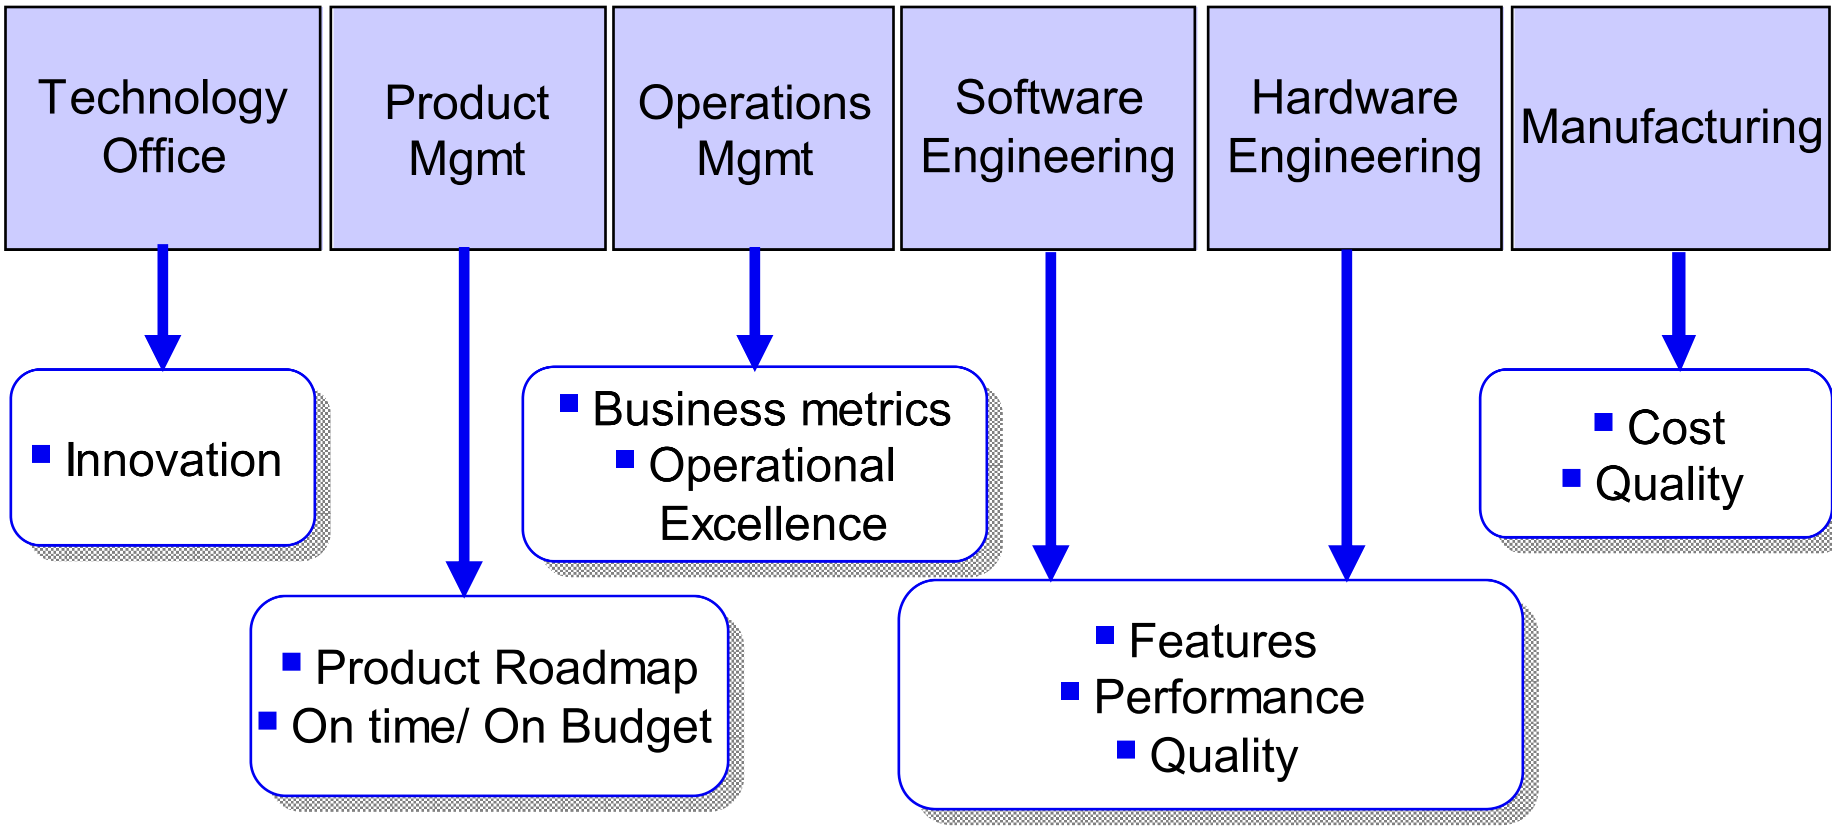

Structured for Performance

Agenda

- Introduction & Overview – Jim Rottsolk
- Sales Strategy – Peter Ungaro
- Purpose Built Systems - Ly Pham
- **Future Directions – Burton Smith**
- Closing Comments – Jim Rottsolk

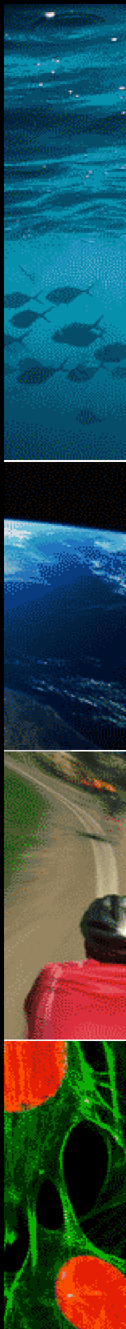

The Supercomputer Company

Future Directions – What's Next

Burton Smith
Chief Scientist

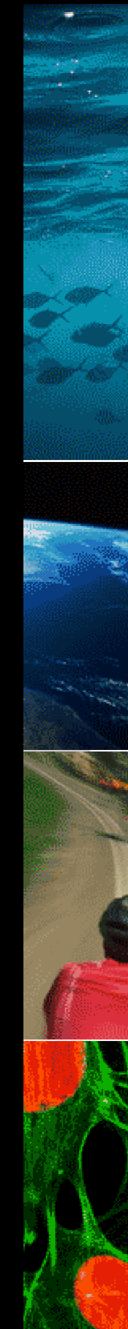

Two current architectural styles

- Sparsely connected, low bandwidth systems
  - Multiple disjoint local address spaces
  - Send and receive
  - PC-intended processors
  - LAN or WAN derived interconnect
- Densely connected, high bandwidth systems
  - Single global address space
  - Load and store anywhere in memory
  - Latency-tolerant processors
  - Custom high bandwidth interconnect
- Cray builds both, and is unique by building the latter
  - What's different about our approach?

Latency, bandwidth, and concurrency

- In a system that transports objects from input to output without creating or destroying them,

  latency x bandwidth = concurrency

- In queueing theory, this result is known as Little's law.

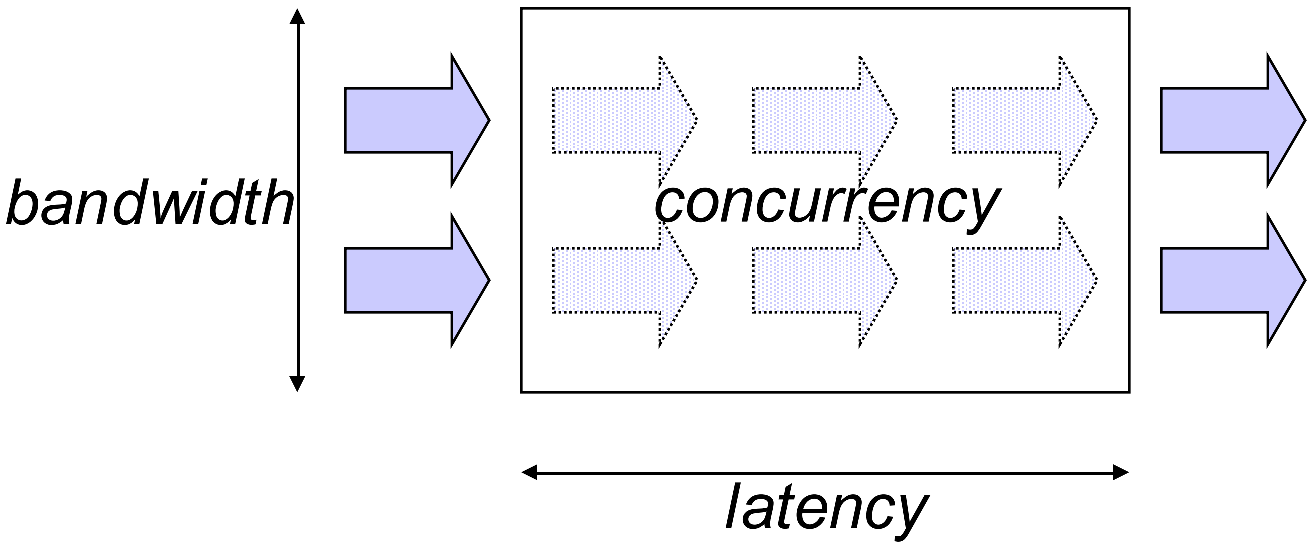

Little's law in action

- Suppose every hour 50,000 people get on airplanes for trips that average 3.5 hours. How many people are in the air?

  concurrency = latency x bandwidth
  = 3.5 hr x 50,000 people/hr
  = 175,000 people

- A microprocessor can only read 16 bytes from its memory at a time, and the memory latency is 0.1 microsecond. What memory bandwidth results?

  bandwidth = concurrency / latency
  = 32 bytes / 0.1 microsecond
  = 320 megabyte/second

Cray scales up high bandwidth systems

- Both vector pipelining and multithreading supply ample concurrency to tolerate large-scale memory latency
- Bandwidth (not latency) limits Cray's practical system size
- Future systems will be differentiated by bandwidth
  - Purchasers will buy the class of bandwidth they need
  - Vendors will make sure their bandwidth scales properly
- A key issue is the cost of bandwidth
  - It dominates cost in big systems
  - How much bandwidth is enough?
  - The answer pretty clearly depends on the system's uses

Measuring bandwidth

- The Linpack Benchmark used for the Top 500 doesn't do it
- The new HPC Challenge benchmarks do a much better job
  - They measure both local and global bandwidth
  - They measure both regular and irregular accesses
  - They include Linpack for completeness
  - They are getting more traction as time passes
- These benchmarks will be discussed at the Top 500 BOF
  - This evening (11/9/04) at SC04 – 5:30 pm in room 304
  - Or look on the web at http://icl.cs.utk.edu/hpcc/

Our customers' needs change quickly

- High Performance Computing, i.e. "How to compute fast," is a rapidly moving target
- The real trick is to predict the future

  "The easiest way to predict the future is to invent it."

  —Alan Kay
- We are constantly inventing the future at Cray:
  - Processors (vectors, multithreading)
  - Interconnect (toroidal meshes, adaptive routing)
  - Programming (auto-parallelizing compilers, co-arrays)
  - Operating systems (checkpoint-restart, job scheduling)
- We have some new ideas we think will be important

Hybrid system ideas

- Mix Mass-Market Microprocessors (MMs) and Cray's latency tolerant processors (LTs) on a shared infrastructure providing:
  - Interconnection network
  - Operating system
  - I/O devices
- Let each kind of processor do what it does best
  - MMs: lower bandwidth work, services, I/O
  - LTs: higher bandwidth work
- Allow communication among MMs and LTs
  - via shared files
  - via message passing
  - via shared memory

New programming language ideas

- Cray is developing a new high level language called *Chapel*
- Our intent is high performance with easy programming
  - Today's high performance languages are hard to use
  - Today's easy-to-use languages have poor performance
- Chapel will be implemented as open source software
  - It will run on any system
  - It will be at its best on high bandwidth platforms

Bandwidth improvement ideas

- Reduce the number and cost of wires
  - Use the highest feasible signaling rate
  - Use better interconnection network topologies
- Use more sophisticated interconnection network routers
  - Transistors get cheaper and faster every year
- Be very careful with data cache architecture
  - Caches can make irregular addressing expensive
  - Cache coherence traffic destroys performance
- Use fiber optics for cabinet-to-cabinet interconnect
  - Once it becomes inexpensive enough

Application ideas

- New HPC applications sometimes need inventing
  - More speed makes entirely new things possible
- Other times, all you need to do is change time-to-solution
  - For example, Nastran on the Cray-1 in 1978
- We are always searching for new applications such as
  - Simulation of crack propagation in solids
  - Direct visualization of complex computer-based designs
  - Data mining and knowledge discovery

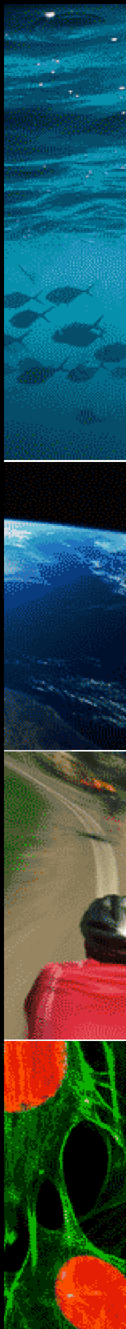

Agenda

- Introduction & Overview – Jim Rottsolk
- Sales Strategy – Peter Ungaro
- Purpose Built Systems - Ly Pham
- Future Directions - Burton Smith
- **Closing Comments – Jim Rottsolk**

Cray Supercomputer Leadership

- Cray is the *only* company dedicated to HPC
  - **Deep experience base**
  - **Uniquely responsive to our customers' needs**
- The Cray approach:
  - **Purpose-built machines for HPC**
  - **Better system balance (flops, memory, network, storage)**
  - **Designed for efficient and reliable scalability**
- Today's solutions:
  - **Cray X1/X1E ⟷ Cray XT3 ⟷ Cray XD1**
- Sustained innovation:
  - ***Rainier:*** **Heterogeneous computing capability in 2006**
  - ***Cascade:*** **Aggressive research program for 2010**

CRAY