Update on Cray Activities in the Earth Sciences

Presented to the 13th ECMWF Workshop on the Use of HPC in Meteorology

3-7 November 2008

Per Nyberg
nyberg@cray.com

Director, Marketing and Business Development

CRAY
THE SUPERCOMPUTER COMPANY

Topics

- General Cray Update
- Cray Technology Directions
  - Overall Technology Focus
  - Specific Updates on:
    - Cray – Intel Collaboration
    - Cray CX1 Announcement
    - ECOphlex Cooling
- Perspectives on Petascale Computing

General Cray Update

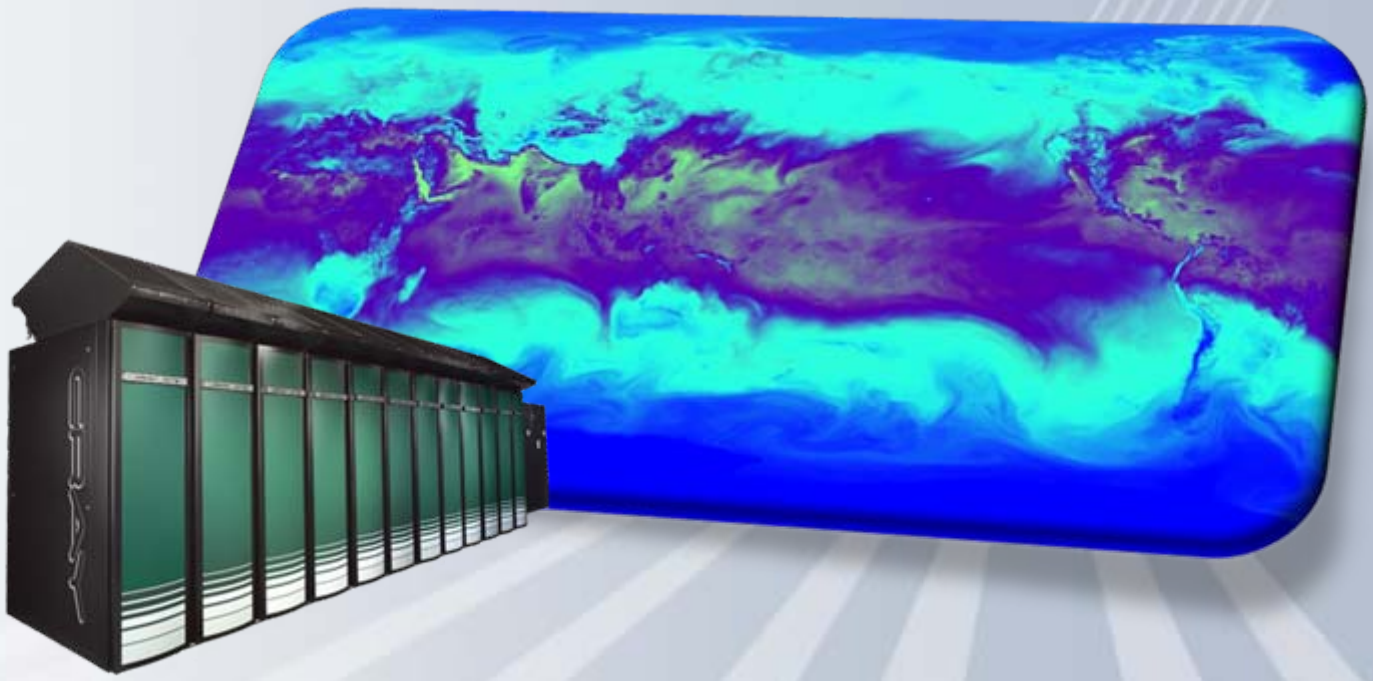

HPC is Cray's Sole Mission

- Significant R&D centered around all aspects of high-end HPC:
  - From scalable operating systems to cooling.
- Cray XT MPP architecture has gained mindshare as prototype for highly scalable computing:
  - From 10's TF to Petaflop.
- First Cray Petaflop supercomputer to be delivered by end of 2008 (ORNL) and second by early 2009 (NSF).

Jaguar Upgrade Brings ORNL Closer to Petascale Computing

- Upgrades to ORNL's Jaguar supercomputer have more than doubled its performance
- The system completed acceptance testing on May 13, running applications in multiple areas such as climate science and astrophysics
- The Jaguar system, a Cray XT4 located at the NCCS, now uses more than 31,000 processing cores to deliver up to 263 trillion calculations a second (or 263 teraflops)

"The Department of Energy's Leadership Computing Facility is putting unprecedented computing power in the hands of leading scientists to enable the next breakthroughs in science and technology."

Thom Mason, ORNL Director

NCCS supercomputer doubles peak performance

The Jaguar XT4

NATIONAL CENTER FOR COMPUTATIONAL SCIENCES

Oak Ridge National Laboratory

ORNL Petascale Jaguar System

- The National Center for Computational Sciences (NCCS) at ORNL is currently installing a new petascale “Jaguar” system.
  - XT5 with ECOphlex liquid cooling.
- Will enable petascale simulations of:
  - Climate science
  - High-temperature superconductors
  - Fusion reaction for the 100-million-degree ITER reactor
- Will be the only open science petascale system in the world when it comes online in 2009.

ORNL Jaguar System

Climate Usage of DoE Cray Systems

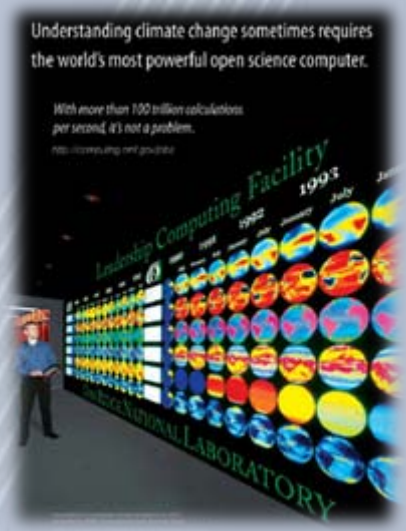

- DoE / NSF Climate End Station (CES)
  - An interagency collaboration of NSF and DOE in developing the Community Climate System Model (CCSM) for IPCC AR5.
  - A collaboration with NASA in carbon data assimilation
  - A collaboration with university partners with expertise in computational climate research.
- DoE / NOAA MoU
  - DoE to provide NOAA/GFDL with millions of CPU hours.
  - Climate change and near real-time high-impact NWP research .
  - Prototyping of advanced high-resolution climate models.

NOAA NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION
UNITED STATES DEPARTMENT OF COMMERCE

Department of Energy to Provide Supercomputing Time to Run NOAA's Climate Change Models

September 8, 2008

The U.S. Department of Energy's (DOE) Office of Science will make available more than 10 million hours of computing time for the U.S. Commerce Department's National Oceanic and Atmospheric Administration (NOAA) to explore advanced climate change models at three of DOE's national laboratories as part of a three-year memorandum of understanding on collaborative climate research signed today by the two agencies.

NOAA Administrator Conrad C. Lautenbacher, Jr. (left) and DOE Under Secretary for Science Dr. Raymond L. Orbach (right).

High resolution (Credit: NOAA)

NOAA will work with climate change models as well as perform near real-time high-impact (non-production) weather prediction research using computing time on DOE Office of Science resources including two of the world's top five most powerful computers – the Argonne National Laboratory's 557 TF IBM Blue Gene/P and Oak Ridge National Laboratory's 263 TF Cray XT4. NOAA researchers will also receive time on DOE's National Energy Research Scientific Computing Center at Lawrence Berkeley National Laboratory.

Advanced, high-resolution climate models from NOAA's Geophysical Fluid Dynamics Laboratory (GFDL) will be prototyped and compared to other models like the NSF-DOE sponsored Community Climate System Model. This partnership is also consistent with the goals of the U.S. Climate Change Science Program, which is responsible for facilitating the creation and application of knowledge of Earth's global environment through research, observations, decision support, and communication. NOAA and DOE scientists play key roles in national and international assessments, for example, the Nobel Prize winning Intergovernmental Panel on Climate Change.

The Climate End Station: Moving to CESM v1.0

CCSM Development — Peter Gent NCAR

Collaborators: ~100 Scientists from DOE Labs, NCAR, and University Community

Development of the fourth version of the CCSM;
One of the biggest improvements is the simulation of ENSO.

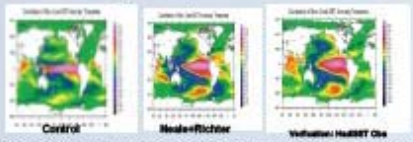

Use of the CCSM to produce short-term climate forecasts to 2030;
Higher atmosphere resolution (0.5º) produces better rainfall in SE

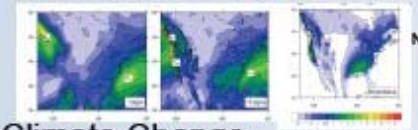

Note improved SE USA rainfall.

Climate Change — Warren Washington NCAR

CCSM Climate Change Working Group: Critical science development and applications in support of climate-informed decision making

- DOE CCSP SAP2.1A *Low Emissions Scenarios
- Coupled ice sheet runs
- Near-Term High-Resolution Climate Prediction
- Climate Change 2100 and beyond
- Prognostic Carbon Aerosol experiments

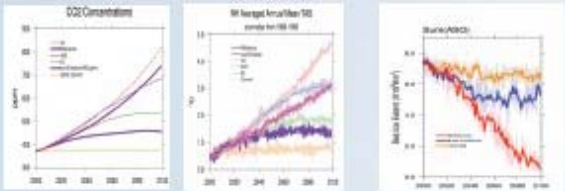

Carbon Assimilation — Dan Anderson NASA

Collaborators: Donald Anderson, Steven Pawson, and Mike Seablom

Simulation and Assimilation of Atmospheric Carbon Species via Multi-year simulations of atmospheric carbon species at 0.5-degree spatial resolution using the GEOS-5 AGCM using specified emission boundary fluxes for CO and CO2. The AGCM will be perturbed, in order to produce estimates of sensitivity of the tropospheric trace gases to different parameters in the model's physical parameterizations. Analysis of the results will quantify how different choices of physical parameters impact the quality of simulations, including agreement with ground-based CO and CO2 measurements and space-based radiances.

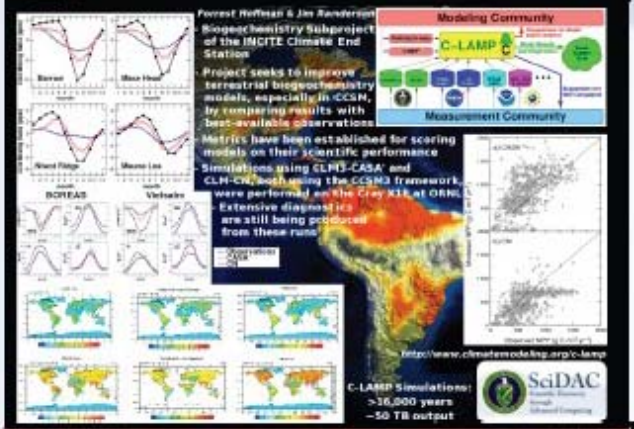

SciDAC ESM w/ Sulfur — Phillip Cameron-Smith LLNL

Status: Components validated. Ocean spin-up underway.

Scientists involved:
Philip Cameron-Smith (PI, LLNL)
Rob Jacob (ANL)
Scott Elliot (LANL)
Phil Jones (LANL)
Mat Maltrud (LANL)
Art Mirin (LLNL)
Jean-Francois Lamarque (NCAR)
John Drake (ORNL)
Dave Erickson (ORNL)
Steve Ghan (PNNL)
+ other SciDAC collaborators
Green = Implementing ESM on Jaguar

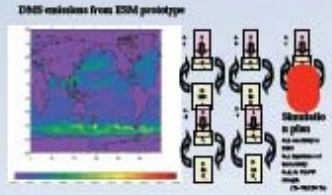

Eddy Resolving POP — Phil Jones LANL

Mathew Maltrud (LANL), Frank Bryan (NCAR), Phil Jones (LANL)

Configure and spin up a 1/10 degree, 42 level stand-alone POP simulation. Then use the result as the initial condition in a fully coupled CCSM run. Compare the results with a similarly configured 1 degree ocean in order to quantify the impact of using an eddying ocean model.

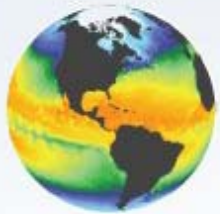

WRF Resolution Studies — Ruby Lueng PNNL

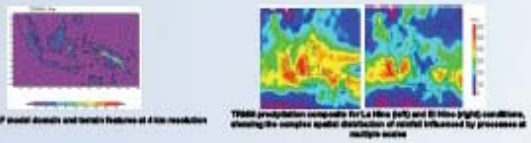

Objective: To systematically assess impacts of spatial resolution on simulating cloud processes and their interactions with the circulation

Approach: Simulations will be performed with the Weather Research and Forecasting (WRF) model One year cloud resolving simulations at 1, 2, and 4 km resolution in the Pacific warm pool. One year mesoscale simulations at 10, 20, and 40 km resolution in the Pacific warm pool. Simulations using the global WRF on an aqua planet with embedded domains at 36 and 12 km resolution in the Pacific warm pool

CAM Development — Jim Hack ORNL

CCSM CAM3 Jan 01 Hour 07

Improvements in mean climate properties were seen when the horizontal resolution in CAM was doubled. However, only modest improvements in transient behavior were observed, suggesting that large-scale systematic biases in mean climate are far more related to the parameterized treatment of non-resolvable physics, such as clouds, convection, radiation and boundary layer. But they also demonstrate that some transient features of the simulation are far more realistic, including the clear presence of tropical storms at the T341 truncation. In addition to the spectral Eulerian dynamical core, we are experimenting with the prototype CAM4 using the FV dynamical core at resolutions as fine as 0.25 dg, which should have solution accuracy equivalent to the T341 spectral model.

NCAR

University Component — David Erickson ORNL

The magnitude and ubiquity of deforestation in the Amazon basin has become increasingly clear in recent years with overwhelming evidence that vegetation is not passively responding to changes in climate and weather, but is dynamically evolving with climate, human and natural disturbances to vegetation. The MIT team (Knox/Bras/Erickson) is simulating this interaction via an ensemble of Amazonian land-atmosphere states via coupled land-atmosphere limited-area model simulations. These simulations will use high and low carbon emission GCM (CCSM3-ORNL) output as lateral boundaries, and forced anthropogenic deforestation lower boundary conditions from both business-as-usual (aggressive deforestation) and conservation type.

WRF 'nature' Benchmark on ORNL Cray XT4

- **Key Performance Enabling Technologies:**
  - Efficient algorithmic implementation in user application.
  - Good scalability at small number of gridpoints per core.
- **Fundamental Cray MPP architectural features for scalability:**
  - MPP 3-D Torus
  - Efficient communications layers
  - Light Weight Kernel / Jitter-free OS

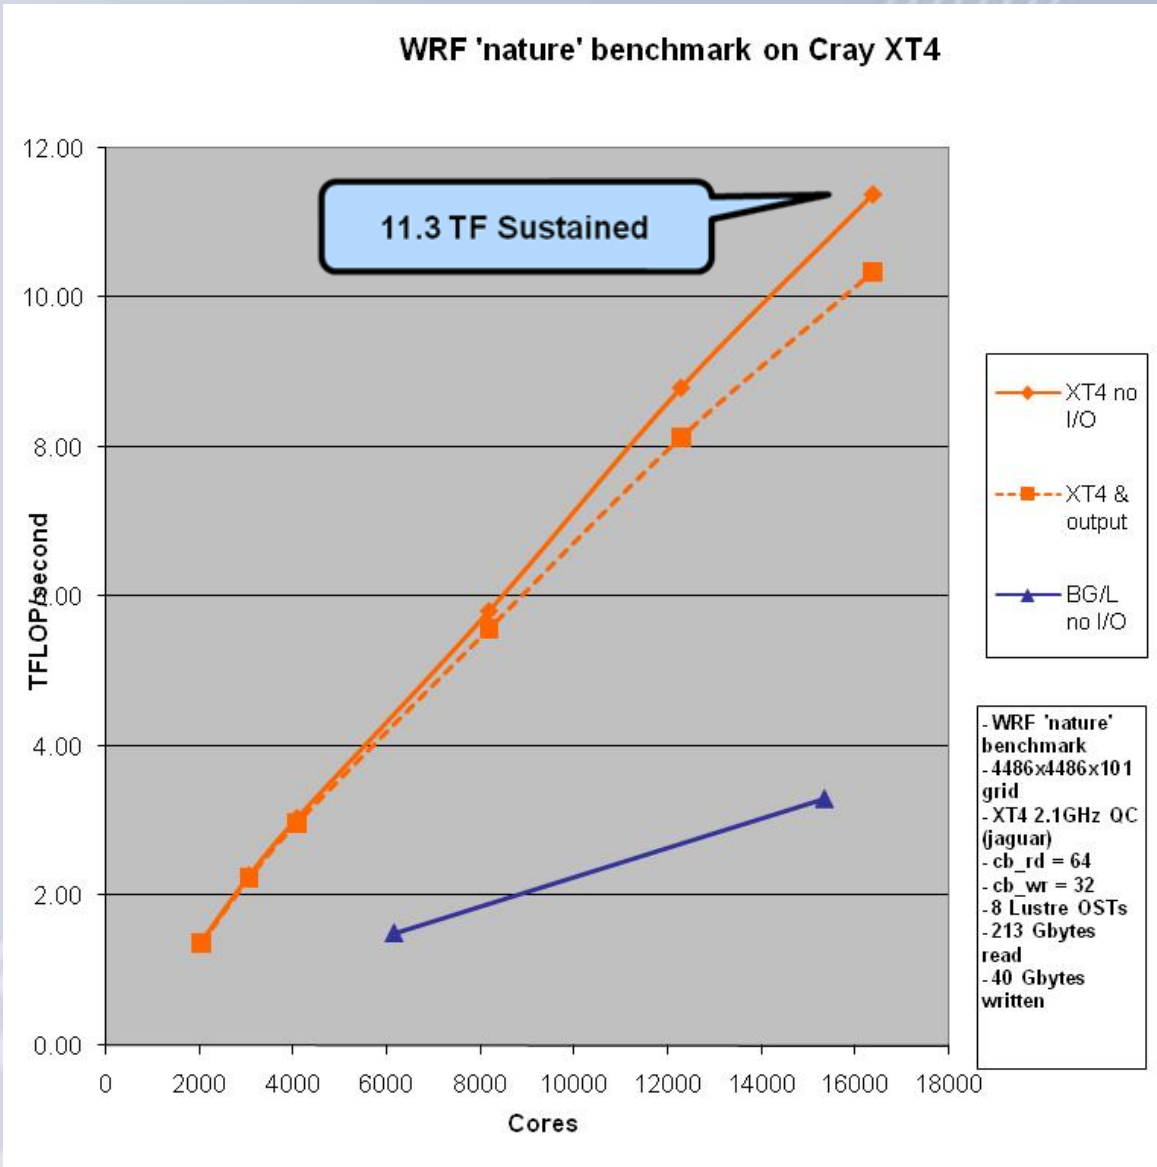

National Science Foundation – Track 2 Award

- Awarded to the University of Tennessee.
- Quad-core AMD Opteron processor system to be installed in 2008.
- Scheduled upgrade to near petascale in 2009.
- Housed at the University of Tennessee – Oak Ridge National Laboratory Joint Institute for Computational Sciences.

THE UNIVERSITY of TENNESSEE
KNOXVILLE, CHATTANOOGA, MARTIN, TULLAHOMA, MEMPHIS

United Kingdom's Engineering and Physical Sciences Research Council (EPSRC)

- Europe's largest HPC initiative - HECToR (High End Computing Terascale Resources)
- Multi-year $85M contract for advanced Cray hybrid system:
  - Cray XT4 + Cray X2
  - Initial peak capability of 60 Tflops
- Installed at the University of Edinburgh's Parallel Computing Centre (EPCC)
- ~20% of the system to be used by NERC (Natural Environment Research Council) in areas such as climate research.

MeteoSwiss and University of Bergen

- In September 2007 CSCS began provision of a Cray XT4 for exclusive operational NWP usage by MeteoSwiss.
- By early 2008 MeteoSwiss was the first European country to run 1~2 km high resolution regional forecasts.

- 52 Tflops XT4 installed in 2008.
- Primary disciplines include marine molecular biology, oceanography and climate research.
- Users include:
  - University of Bergen
  - Institute of Marine Research
  - Nansen Environmental and Remote Sensing Center
  - NOTUR

Danish Meteorological Institute

- DMI win was announced in November 2007.
- Will significantly enhance DMI's NWP and climate assessment capabilities.
- Dual XT5 system configuration for operations and research with full failover.
- Current operations has transitioned from previous platform to running on single cabinet XT4.
- Will transition next operational suite to XT5s in early 2009.

Cray Technology Directions

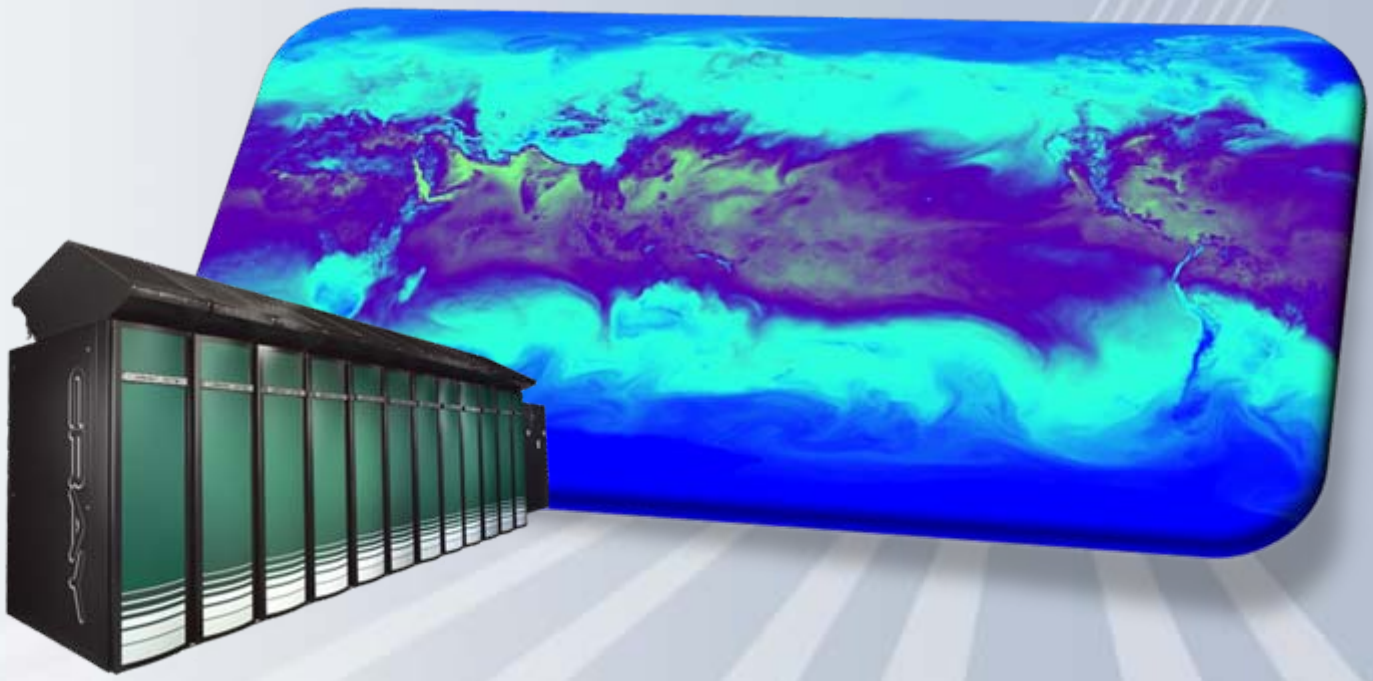

Opposing Forces in Today's HPC Marketplace

April 25, 2008

HPC Innovation in the Era of 'Good Enough'

by John E. West , for HPCwire

Supercomputing is increasingly about managing massive parallelism:

- **Key to overall system performance, not only individual application performance.**
- **Demands investment in technologies to effectively harness parallelism (hardware, software, I/O, management, etc…).**

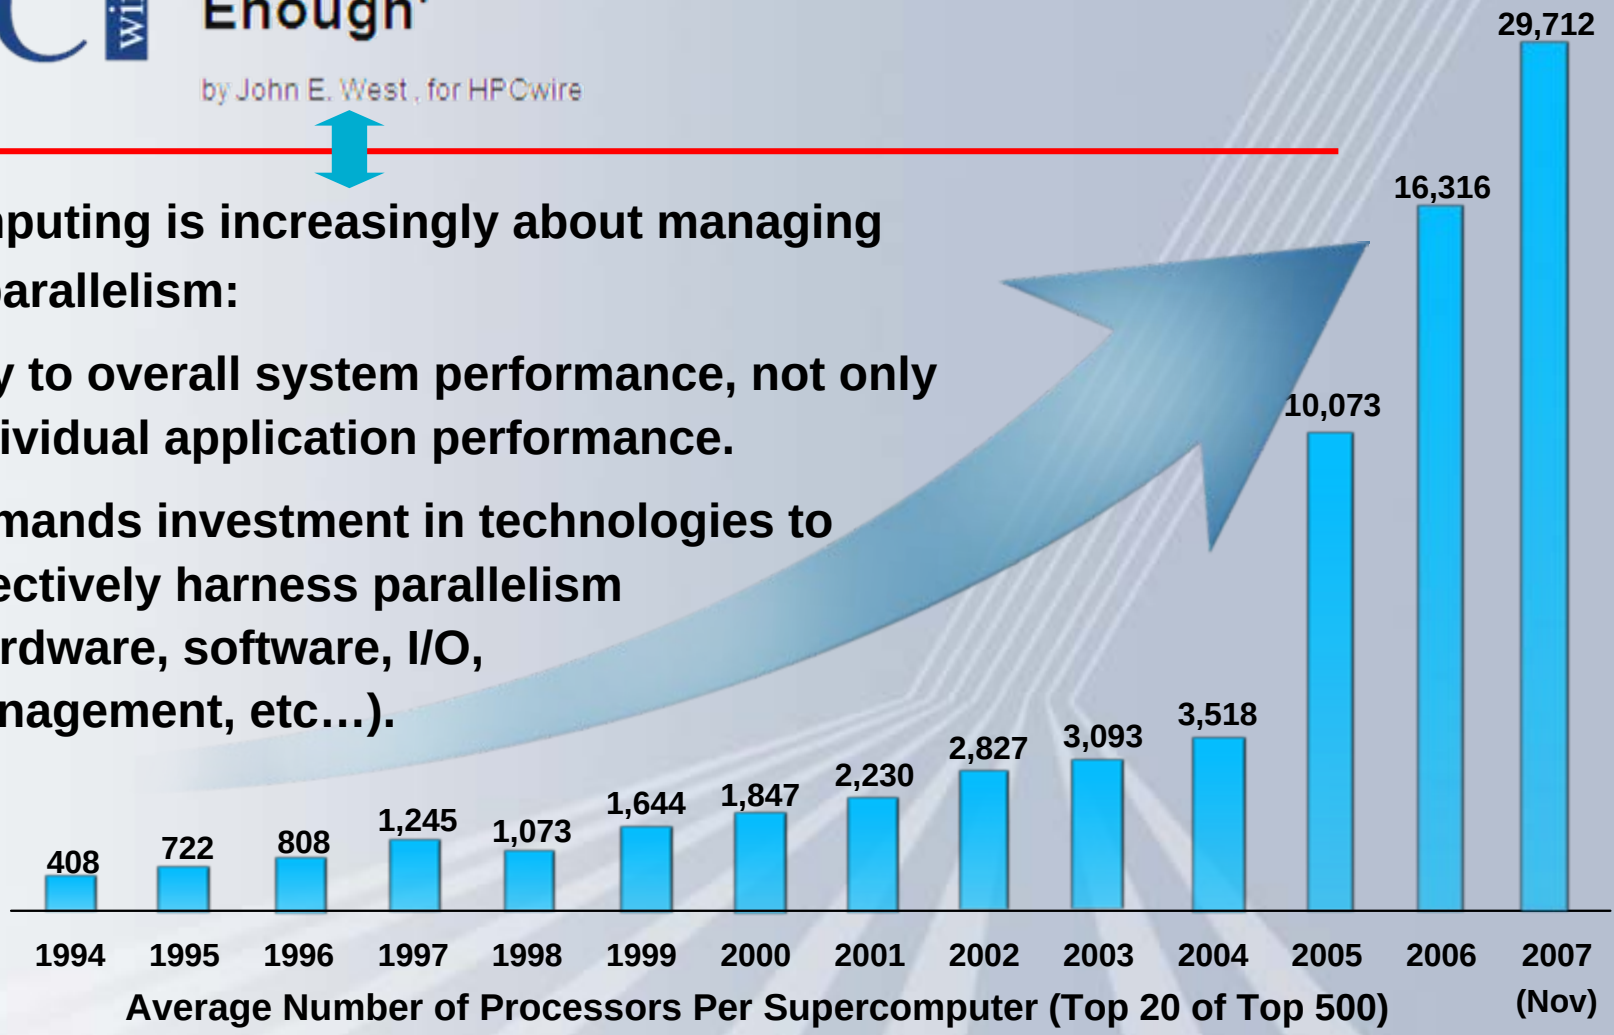

Source: www.top500.org

The Cray Roadmap

Realizing Our Adaptive Supercomputing Vision

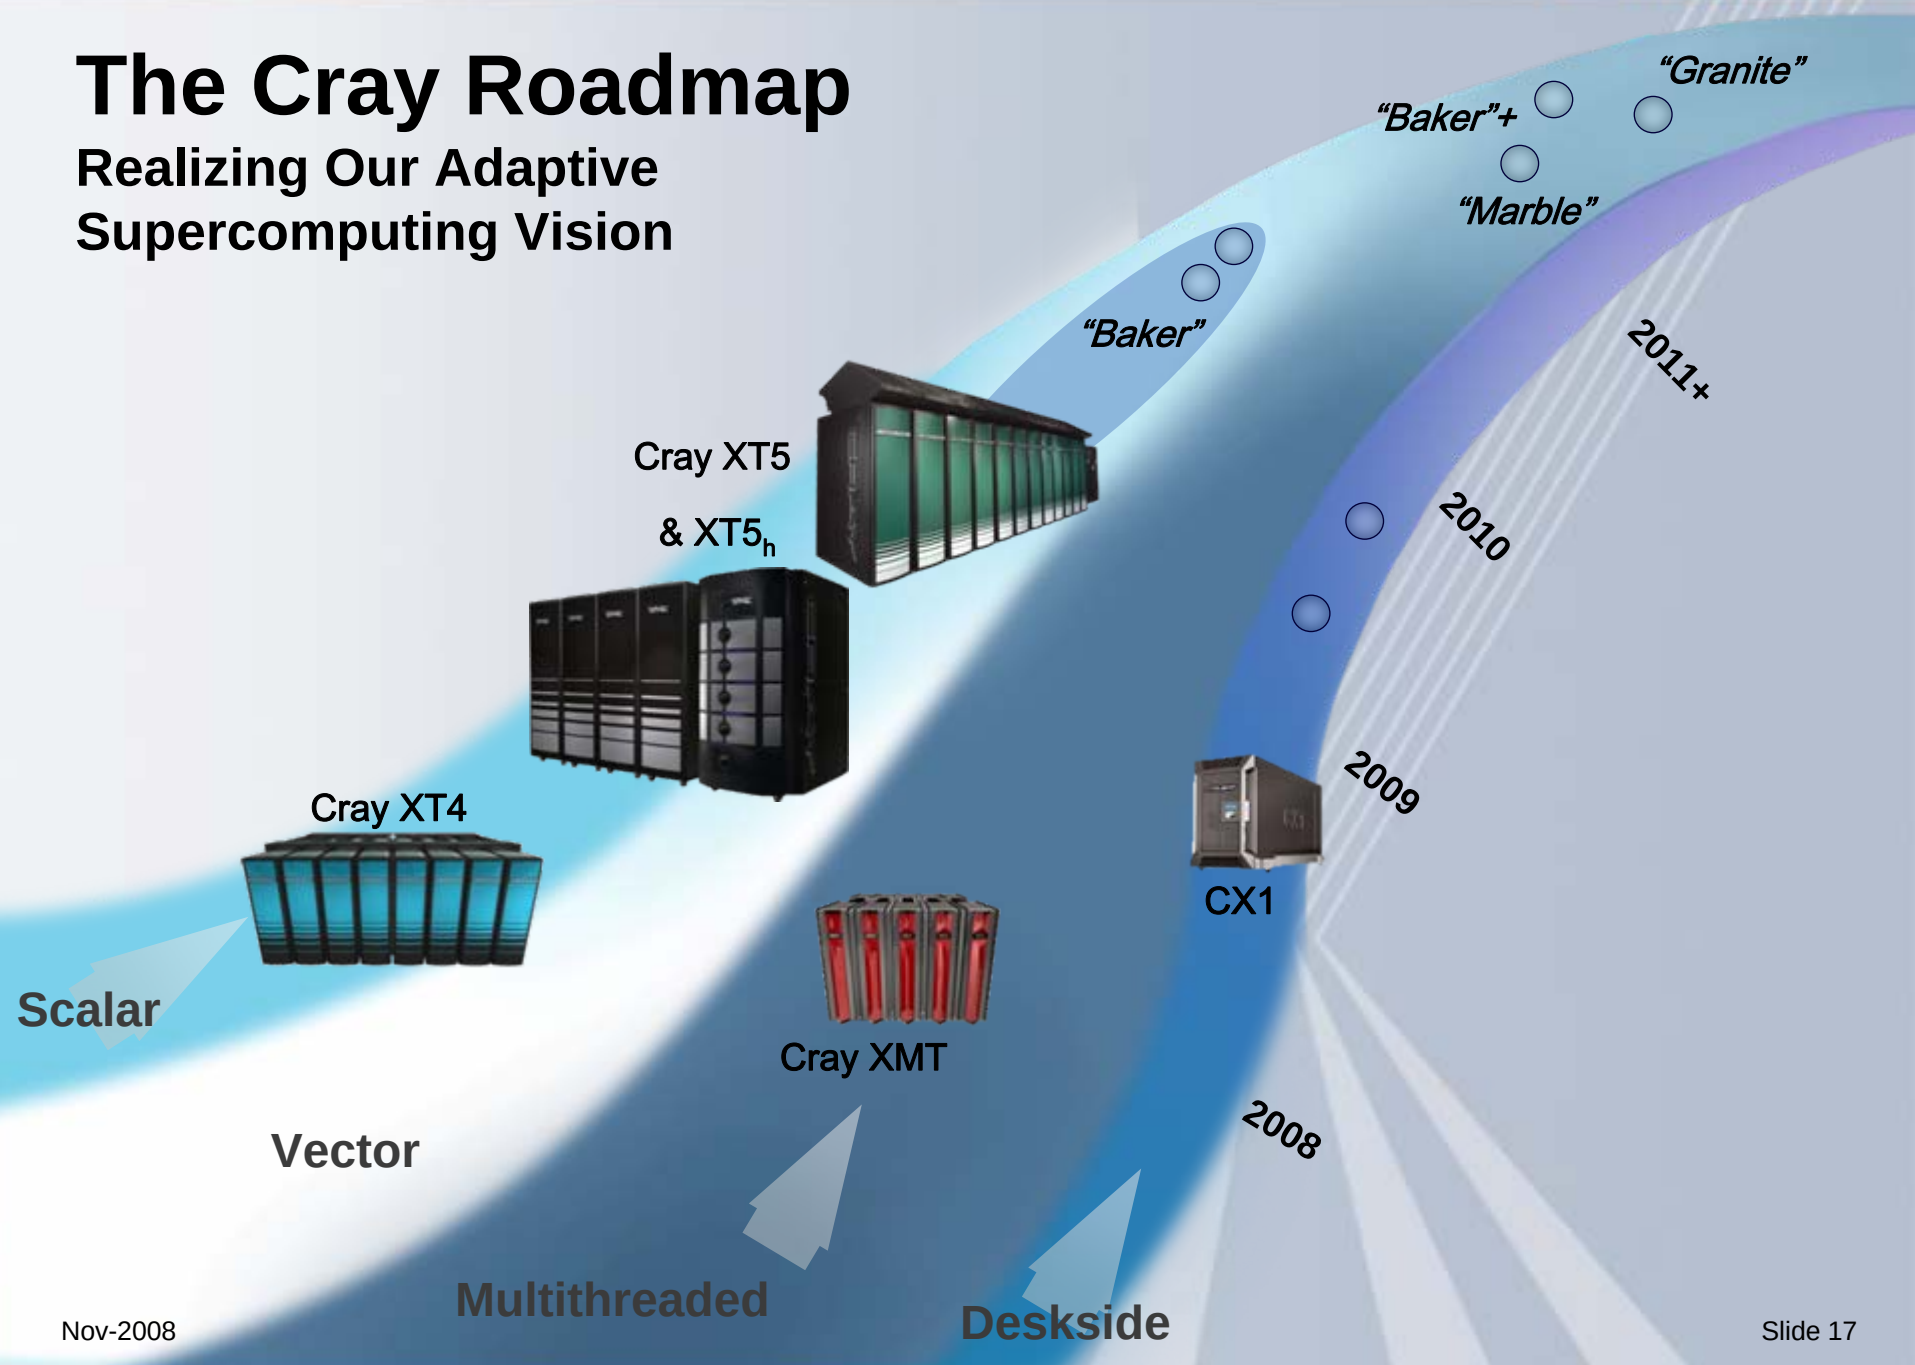

Cray Technology Directions

- Scalable System Performance:
  - Enabling scalability through technologies to manage and harness parallelism.
  - In particular interconnect and O/S technologies.
    - Cray is a pioneer and leader in light weight kernel design.
- Green Computing: packaging, power consumption and ECOphlex cooling:
  - Simple expansion and in-cabinet upgrade to latest technologies, preserving investment in existing infrastructure.
  - Lowest possible total electrical footprint.
- Provide investment protection and growth path through R&D into future technologies: HPCS Cascade Program.

Cray and Intel to Collaborate to Develop Future Supercomputing Technologies

- Announced on April 28, 2008.
- Comprehensive multi-year agreement to advance high-performance computing on Intel processors while delivering new Intel and Cray technologies in future Cray systems.
- Greater choice for Cray customers in processor technologies.
- Cray and Intel will explore future supercomputer component designs such as multi-core processing and advanced interconnects.
- First HPC products will be available in the Cascade timeframe.

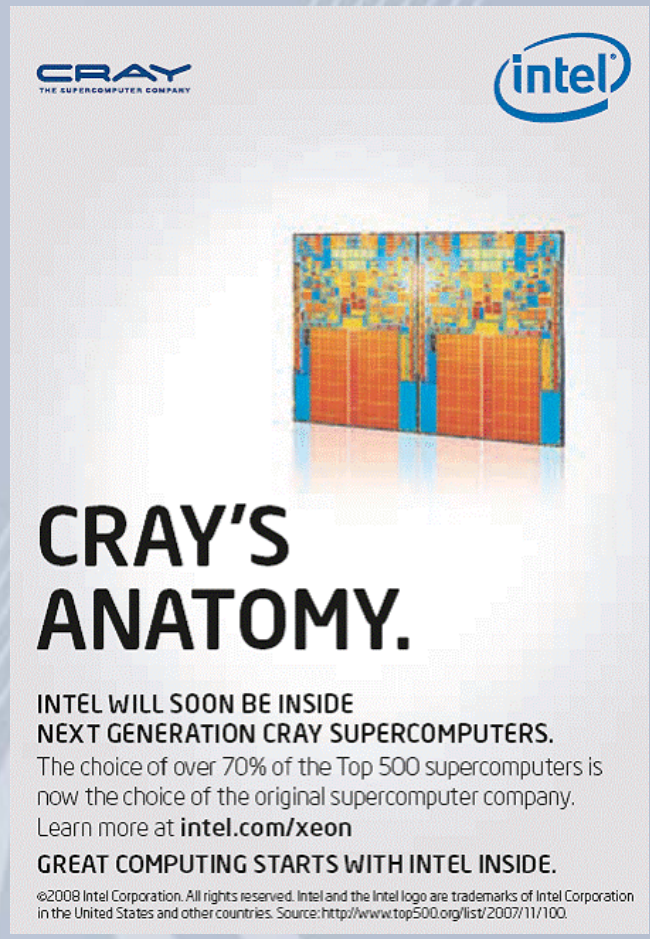

Cray CX1 Deskside Supercomputer

- **Certified Cray Quality**
- **Best in Class Customer Experience**
- **Integrated HPC Operating Environment**

Windows HPC Server 2008

redhat

- **World's Most Versatile Cluster**
  - Single or dual socket compute node
  - High-end graphic node
  - High-capacity storage node
  - Up to 4TB of internal storage
  - Built-in Gigabit or InfiniBand Interconnect
- **Office and lab compatible with standard office power outlet**

intel Xeon inside

We Take Supercomputing Personally

Two Approaches to Cooling – Single Infrastructure High Efficiency Cabinet can be Liquid or Air-Cooled

- Liquid Cooled – HE Cabinet with ECOphlex Technology
  - No change in ambient room temperature.
  - Minimizes power consumption by reducing need for air handler/conditioners.
  - Elimination of Chillers can be the single biggest step toward “Green”
  - Critical for situations where little or no additional air cooling is available.
- Air Cooled – HE Cabinet
  - Utilizing each cubic foot of cooling air to the extreme
  - Minimal total volume of cooling air required.
  - Minimal fan electrical power consumed as compared to muffin fans.
- Made possible by
  - Proprietary Axial Turbine Blower
  - Progressive heat sink technology blade design
  - Progressive air cooled chassis design.

PHase change Liquid EXchange(PHLEX)

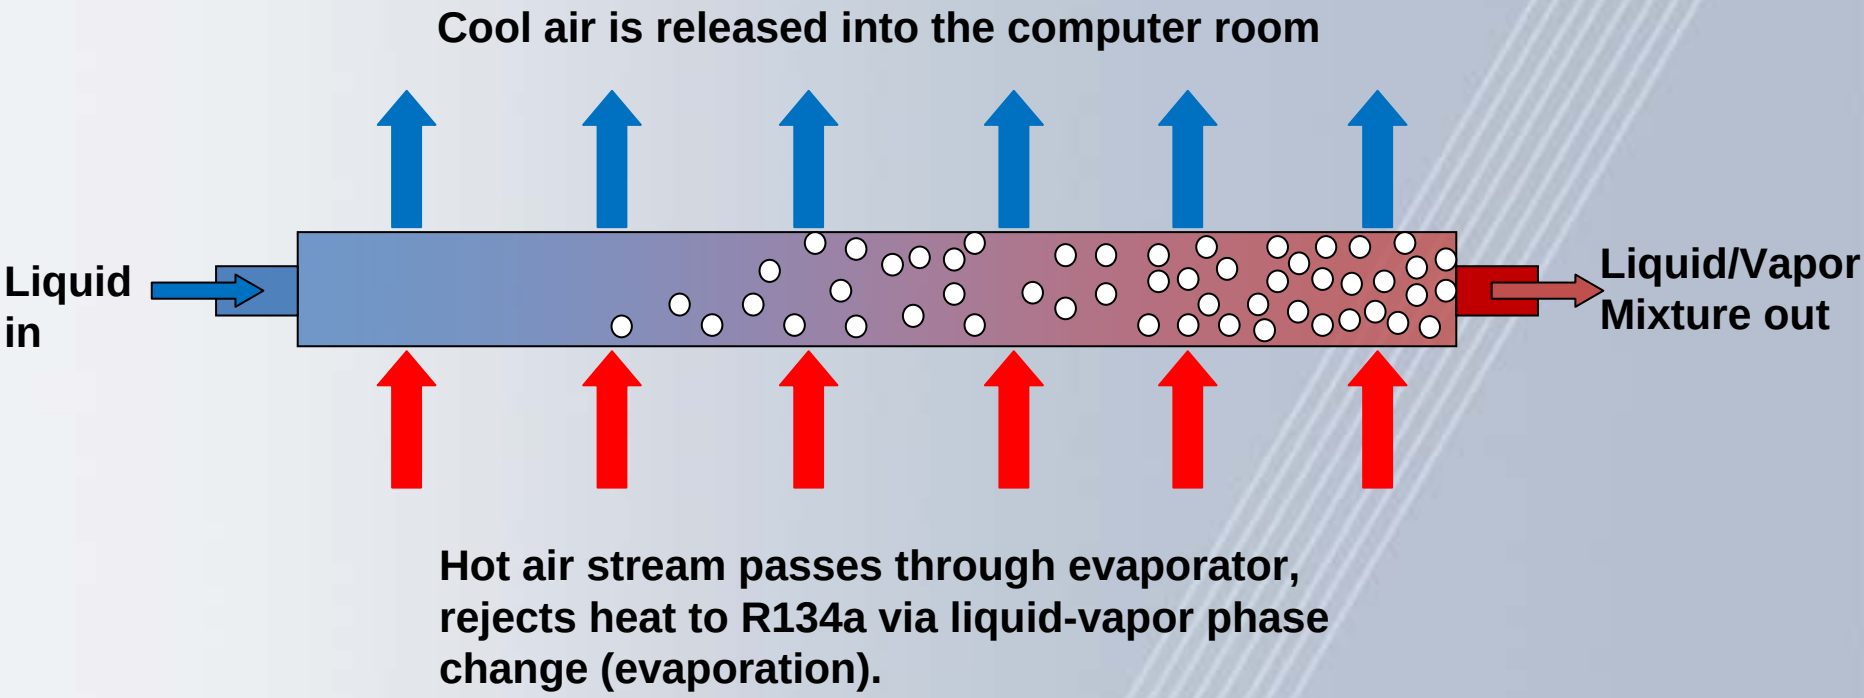

- R134a absorbs energy only in the presence of heated air.
- Phase change is 10x more efficient than pure water cooling.
- *Corollary*: Weight of coils, fluid, etc. is 10X less than water cooling

XT5 System with ECOphlex

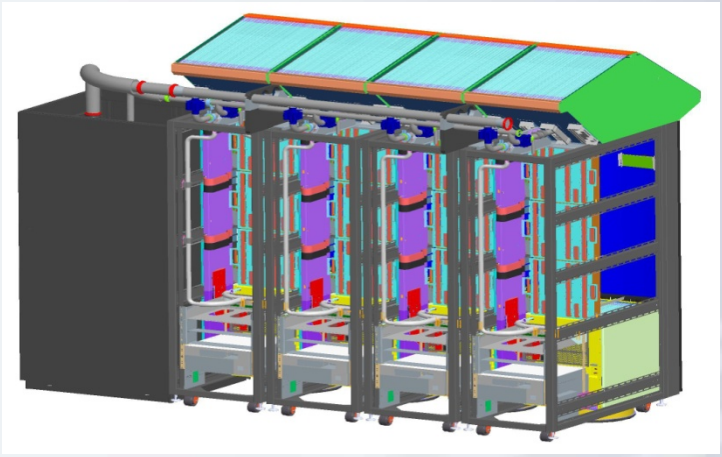

Advantages of ECOphlex vs. Water-Cooled Solutions

- No water near computer components
  - If leaked, will not damage components
  - No condensate drains or chance of water damage
- Lightweight
  - Small volume of coolant in the compute cabinet.
  - Floor load is similar to air-cooled cabinets
- Cost of Ownership
  - Coolant can be re-condensed using building water.
  - In many cases, cooling can be “free”
- Serviceability and Reliability
  - Blades are still air cooled and can be pulled while system in operation.
  - Large systems can continue to fully function if an HEU is down.

Perspectives on Petascale Computing

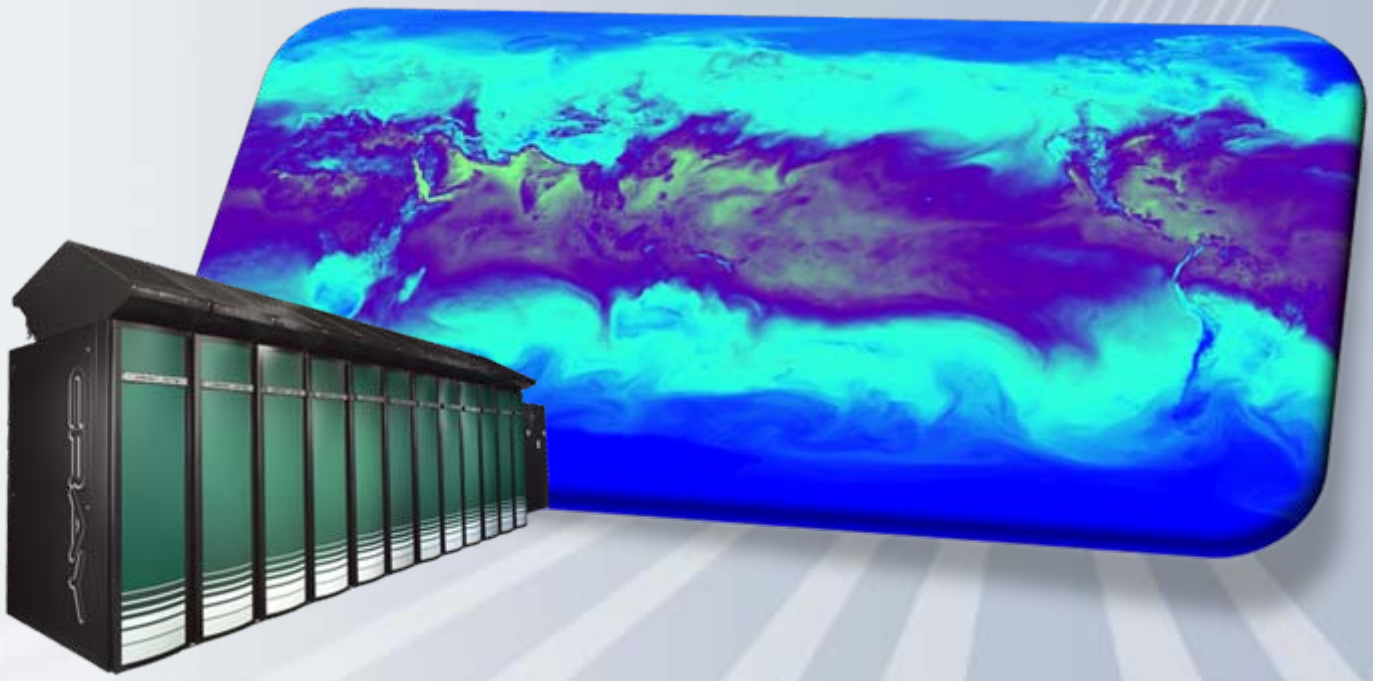

Petascale (or even Terascale) Perspectives

- There remains a tremendous number of models and application areas that have not yet reached even the terascale level.
- Those that can scale have benefited from a focused, iterative multi-year algorithmic optimization effort:
- Optimization strategies do not remain stagnant and must take advantage of evolving hardware and software technologies.
- Ongoing access to scalable, leadership class systems and support is essential:
  - Either you have actually run on 5K, 10K, 20K, 50K… processors, or you have not ! ***Theory vs. Reality***

Evolving Path to Petascale

- Three basic paths are available to users today:
  - Standards based MPP
  - Low power
  - Accelerators
- Cray Cascade program is intended to bring these together to provide greater flexibility within a single architecture.
- Will be accomplished through:
  - Continued investment in core MPP and emerging technologies.
  - Continued investment in customer relationships.
- Cray's wealth of experience in pushing the boundaries of scalability will continue to positively impact entire HPC community.


Thank you for
your attention.
CRAY
THE SUPERCOMPUTER COMPANY