Trends in Supercomputing 2011-2020
The Road to Exascale

April 2011

Barry Bolding, Ph.D.
VP, Cray Product Division

Adenda

- Corporate Overview
- Trends in Supercomputing
  - Types of Supercomputing and Cray's Approach
  - The Cloud
  - The Exascale Challenge
- Conclusion

Corporate Overview

Cray Today…

Seymour Cray founded Cray Research in 1972

- SGI purchased Cray Research in 1996

Cray Inc. formed April 2000

- Tera purchased Cray Research assets from SGI
- Nasdaq: CRAY
- 850 employees across 20 countries
- Headquarters in Seattle, WA

Four Major Development Sites:

- Chippewa Falls, WI
- St. Paul, MN
- Seattle, WA
- Austin, TX

HPCwire 2009 Readers' And Editors' Choice Awards

Our Goals

Market Leadership

- High-end of HPC
- $2.0B Market Opportunity
- Leverage Scalability Advantages
- Goal: #1 in Capability Segment

Sustained Profitability

- Grow Top-Line
- Improve Product GM's
- Reduce OPEX (as % of Rev)
- Drive Positive Bottom-Line

Focus on Innovation and Execution

Cray Technology Innovations

CRAY THE SUPERCOMPUTER COMPANY

System Interconnect	Custom interconnect and communications network
Systems Management & Performance	Software to productively manage and extract performance out of thousands of processors as a single system
Packaging	Very high density, upgradeability, liquid and air-cooling
Adaptive Supercomputing	Single integrated system

Building the Technologies and Infrastructure for Superior Scalability and Sustained Performance

PNNL and Supercomputing

- PNNL has a long history in Supercomputing
  - Long tradition working with Cray, IBM, HP
- EMSL – Environmental Molecular Sciences Lab
  - Software and Science: NWCHEM, Global Arrays, ARMCI
- CASS-MT – Center for Adaptive Supercomputing Software
  - Moe Khaleel, John Feo
    - Collaboration between PNNL,Cray,Georgia Tech, Sandia National Labs
    - Cray XMT, 128 Multithreading processors, 1TB memory
  - Social Network Analysis
  - Analysis of Power Grids

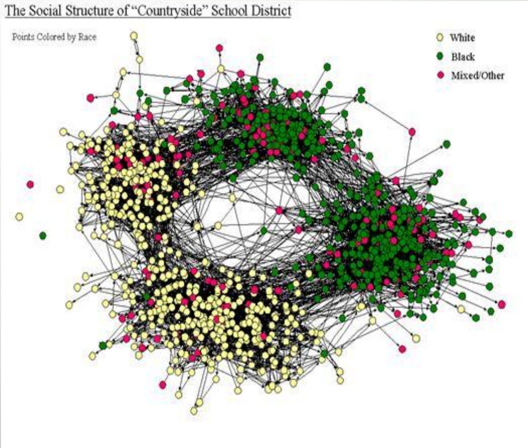

Trends in Supercomputing

Cray Is Supercomputing

We build the world's largest and fastest supercomputers for the highest end of the HPC market

Targeting government agencies, research institutions and large enterprises

We help solve the "Grand Challenges" in science and engineering that require supercomputing

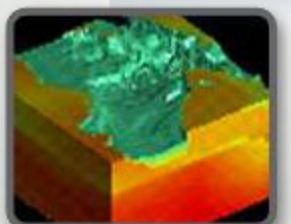

Earth Sciences
EARTHQUAKE PREDICTION

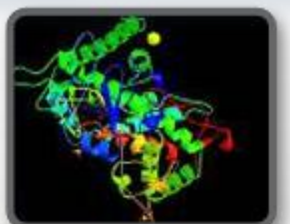

Life Sciences
PERSONALIZED MEDICINE

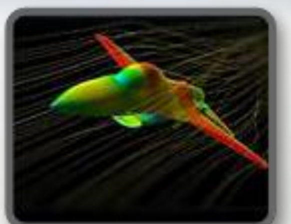

Defense
AIRCRAFT DESIGN

Computer-Aided Engineering
CRASH SIMULATION

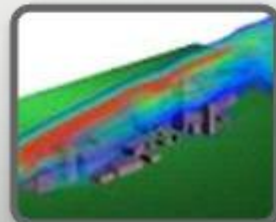

National Security
THREAT ANALYSIS

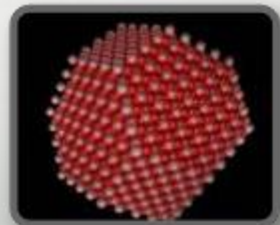

Scientific Research
NANOFUEL DEVELOPMENT

Computing over Mathematics – Superconducting Materials

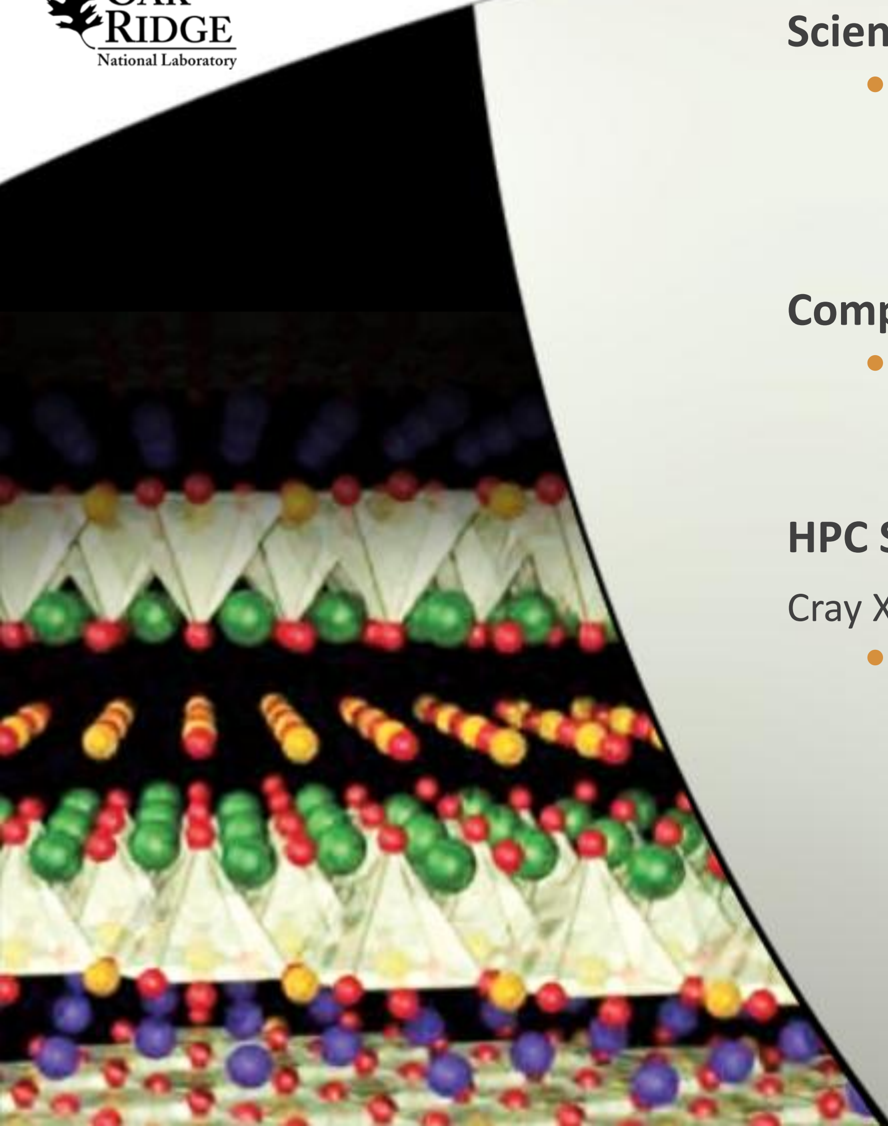

Science Challenge

- Find a superconductor that will exhibit its desirable characteristics – strong magnetic properties and the ability to conduct electricity without resistance or energy loss – without artificial cooling

Computational Challenge:

- Study chemical disorder in high temperature superconductors and the repulsion between electrons on the same atom

HPC Solution

Cray XT5™ supercomputer "Jaguar"

- Modified the algorithms and software design of its DCA++ code to maximize speed without sacrificing accuracy, achieving 1.352 petaflops and the first simulations with enough computing power to move beyond perfectly ordered materials

Understanding superconductors may lead to saving significant amounts of energy

Fastest Systems in the World on Real Applications

CRAY
THE SUPERCOMPUTER COMPANY

OAK RIDGE
National Laboratory

Cray XT5 (Jaguar)
2.3 Petaflops

NERSC

Cray XE6 (Hopper)
1.3 Petaflops

"Supercomputing modeling and simulation are changing the face of science and sharpening America's competitive edge."

Secretary Steven Chu
U. S. Department of Energy

TOP500®
SUPERCOMPUTER SITES

CRAY
THE SUPERCOMPUTER COMPANY

CRAY
THE SUPERCOMPUTER COMPANY

Scalable Performance

Production Efficiency

Adaptive Supercomputing

Cray Software Ecosystem

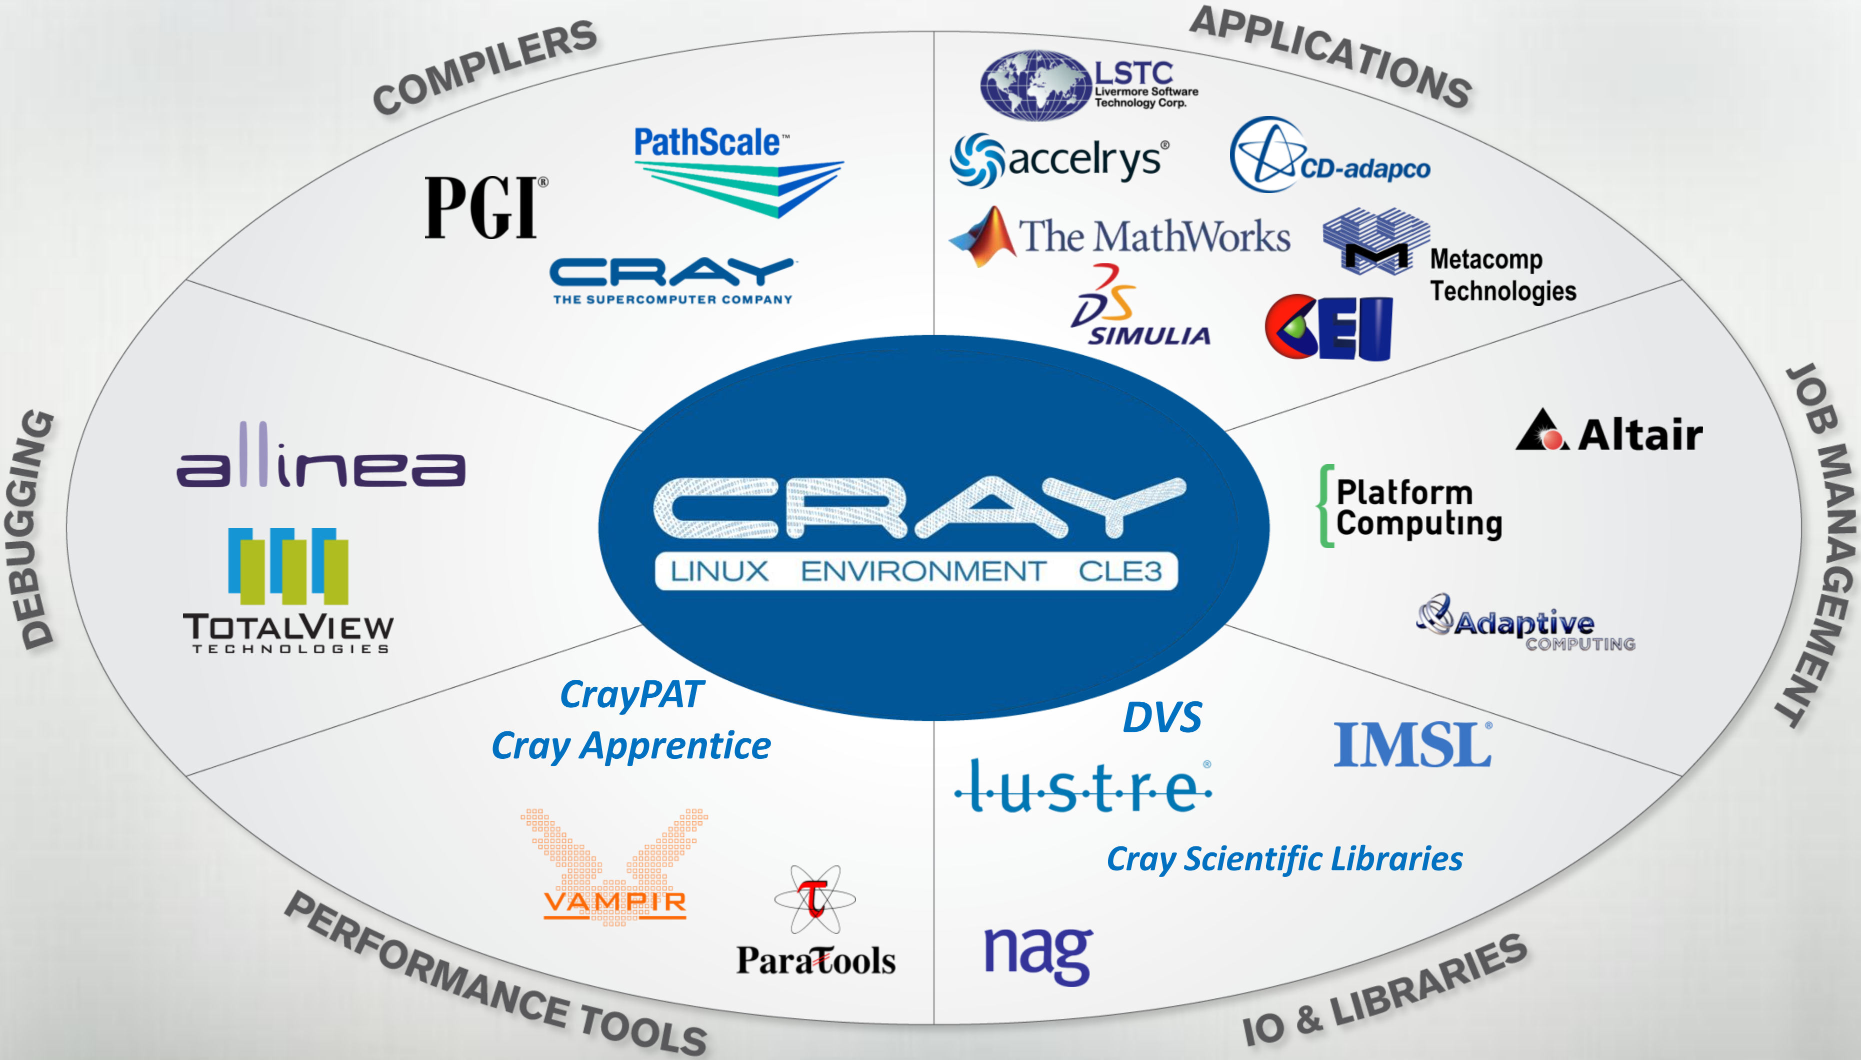

Computing over Data

Power Distribution Networks

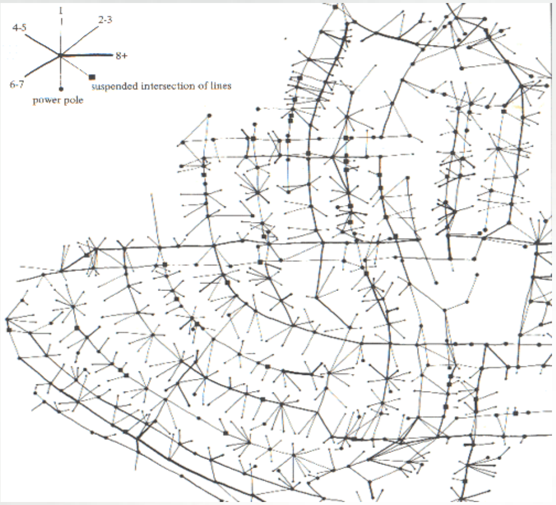

Internet backbone

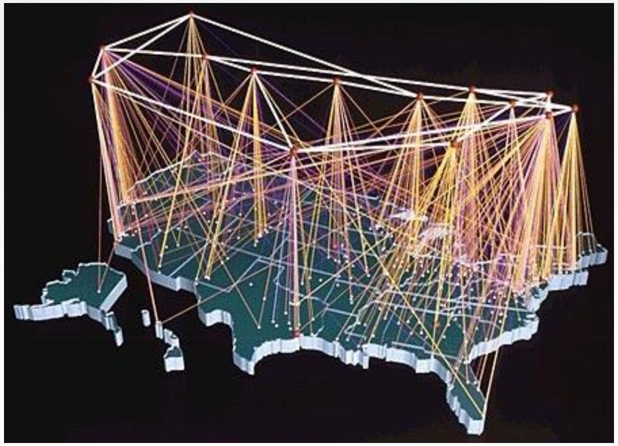

Social Networks

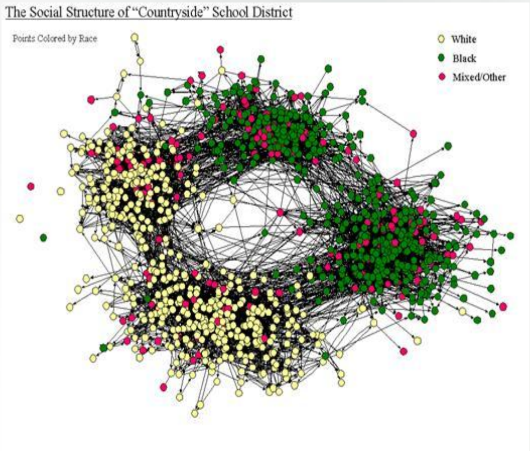

Graphs are everywhere!

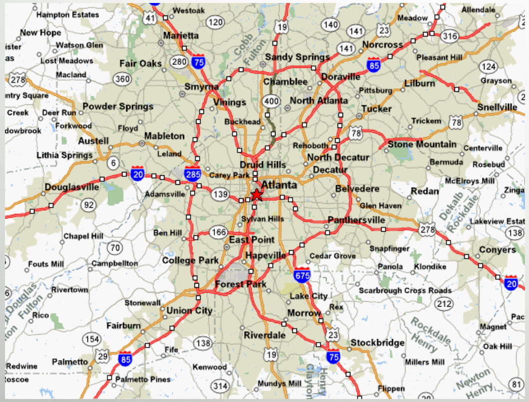

Ground Transportation

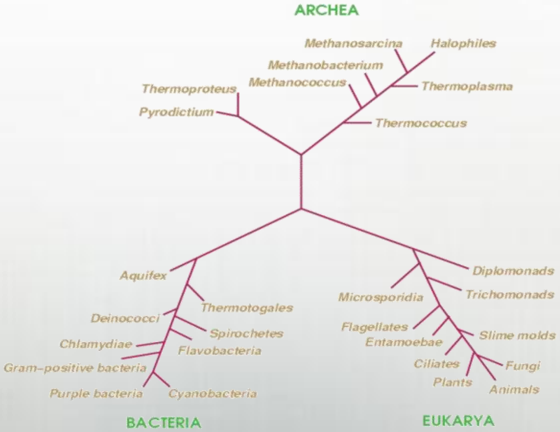

Tree of Life

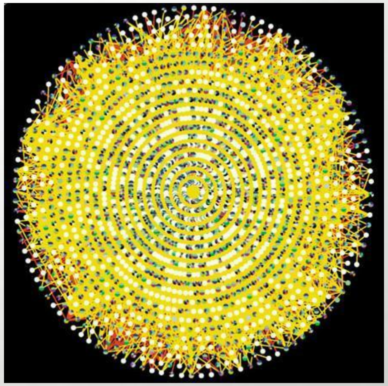

Protein-interaction networks

Custom Engineering Practices

Knowledge Management

Cray XMT Massively Multithreaded technology designed for large scale data analysis on unstructured data

Special Purpose Devices

Custom hardware and software designed to meet critical customer requirements

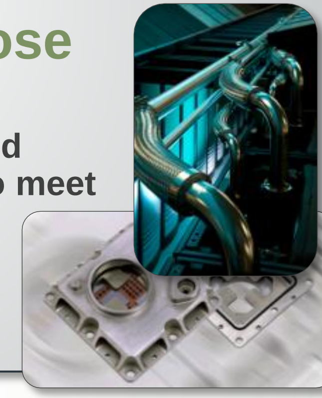

Differentiated, Technology-Led Professional Services

Data Management

Leading data management and storage technologies to externalize key services from supercomputers to data centers

XMT – For Computing over Data

CRAY THE SUPERCOMPUTER COMPANY

XMT Design Goals:

- Increased Memory Capacity
- Improved Reliability, Availability, Serviceability (RAS)
- Reduced Footprint per TB – Power and Space
- Improved Price/Performance
- System and User Software Improvements

Future XMT Systems

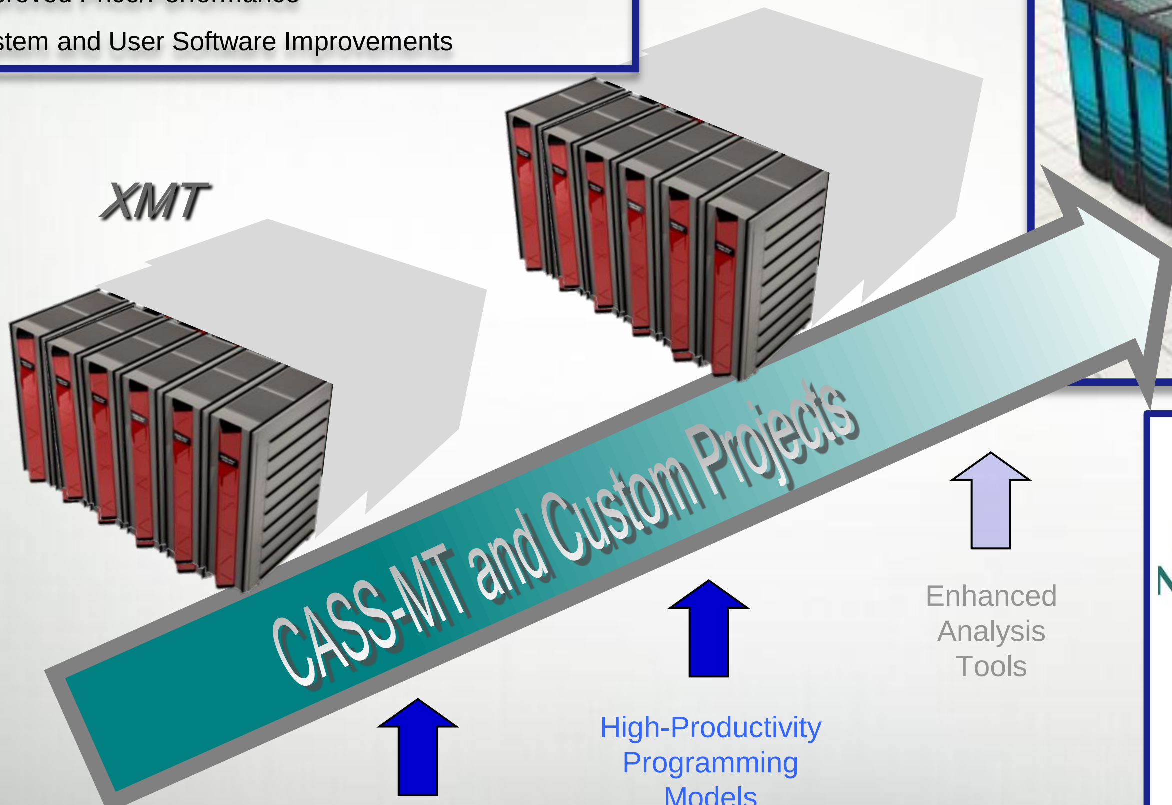

CRAY THE SUPERCOMPUTER COMPANY

Pacific Northwest National Laboratory
Operated by Battelle for the U.S. Department of Energy

Sandia National Laboratories

Washington State University
World Class. Face to Face.

University of Delaware

The Department of Defense

Georgia Tech

The Cloud

The Cloud – Definition

- *On-demand self-service.* Unilateral provisioning
- *Ubiquitous network access.*
- *Location independent resource pooling.* The customer generally has no control or knowledge over the exact location of the provided resources.
- *Rapid elasticity.* Capabilities can be rapidly and elastically provisioned to quickly scale up and rapidly released to quickly scale down.
- *Pay per use.*

- http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.pdf

The Cloud – Definition

- *On-demand self-service.* Unilateral provisioning

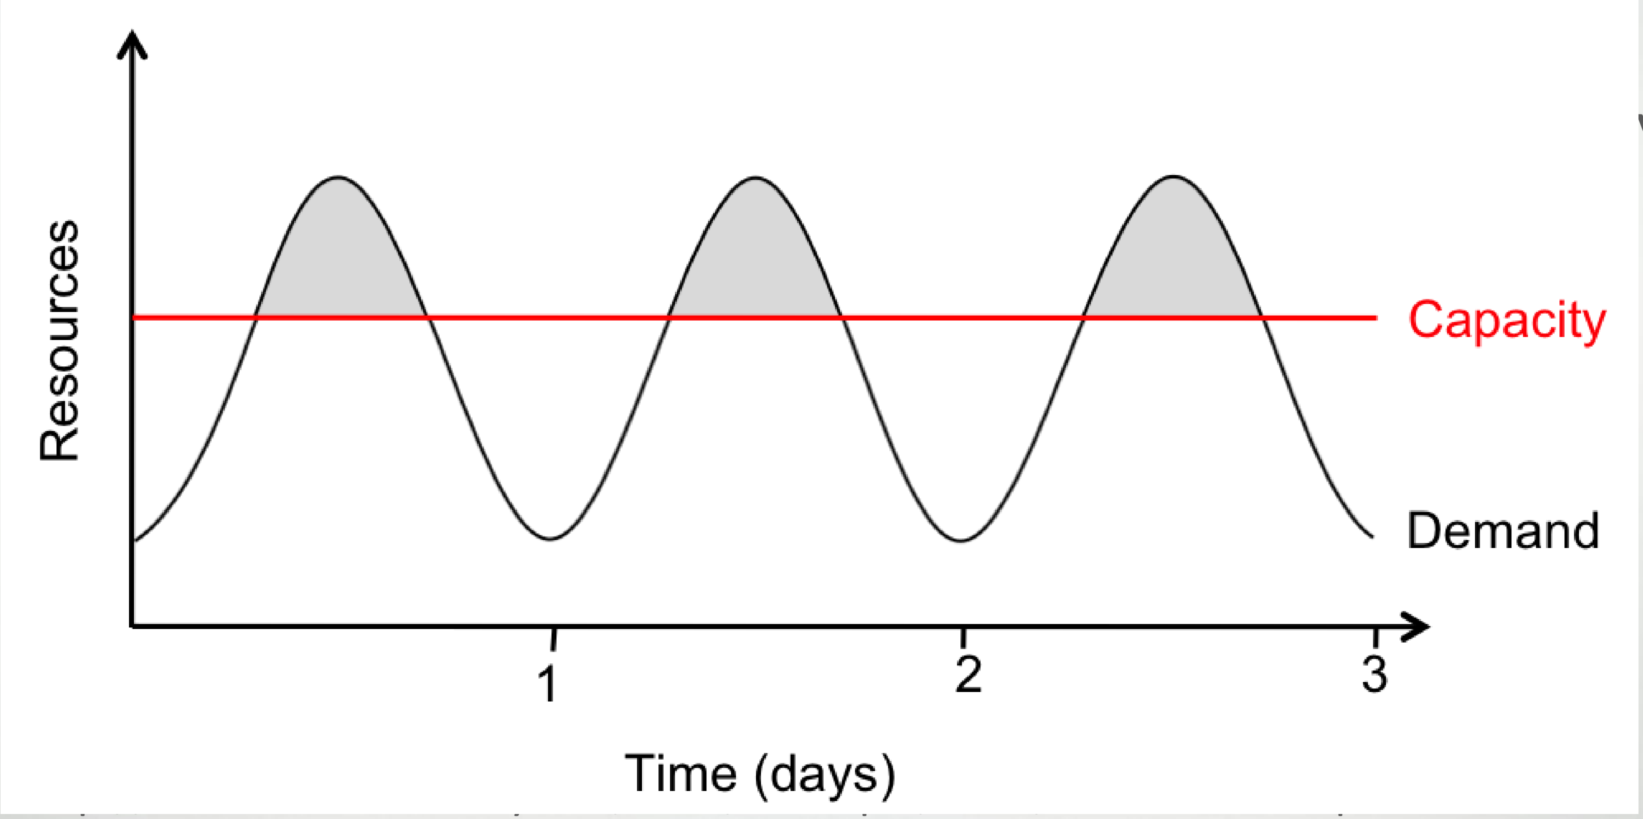

Delivery Models

- *Cloud Software as a Service (SaaS).* The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure
- *Cloud Platform as a Service (PaaS).* The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created applications
- *Cloud Infrastructure as a Service (IaaS).* The capability provided to the consumer is to rent processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software

Delivery Models

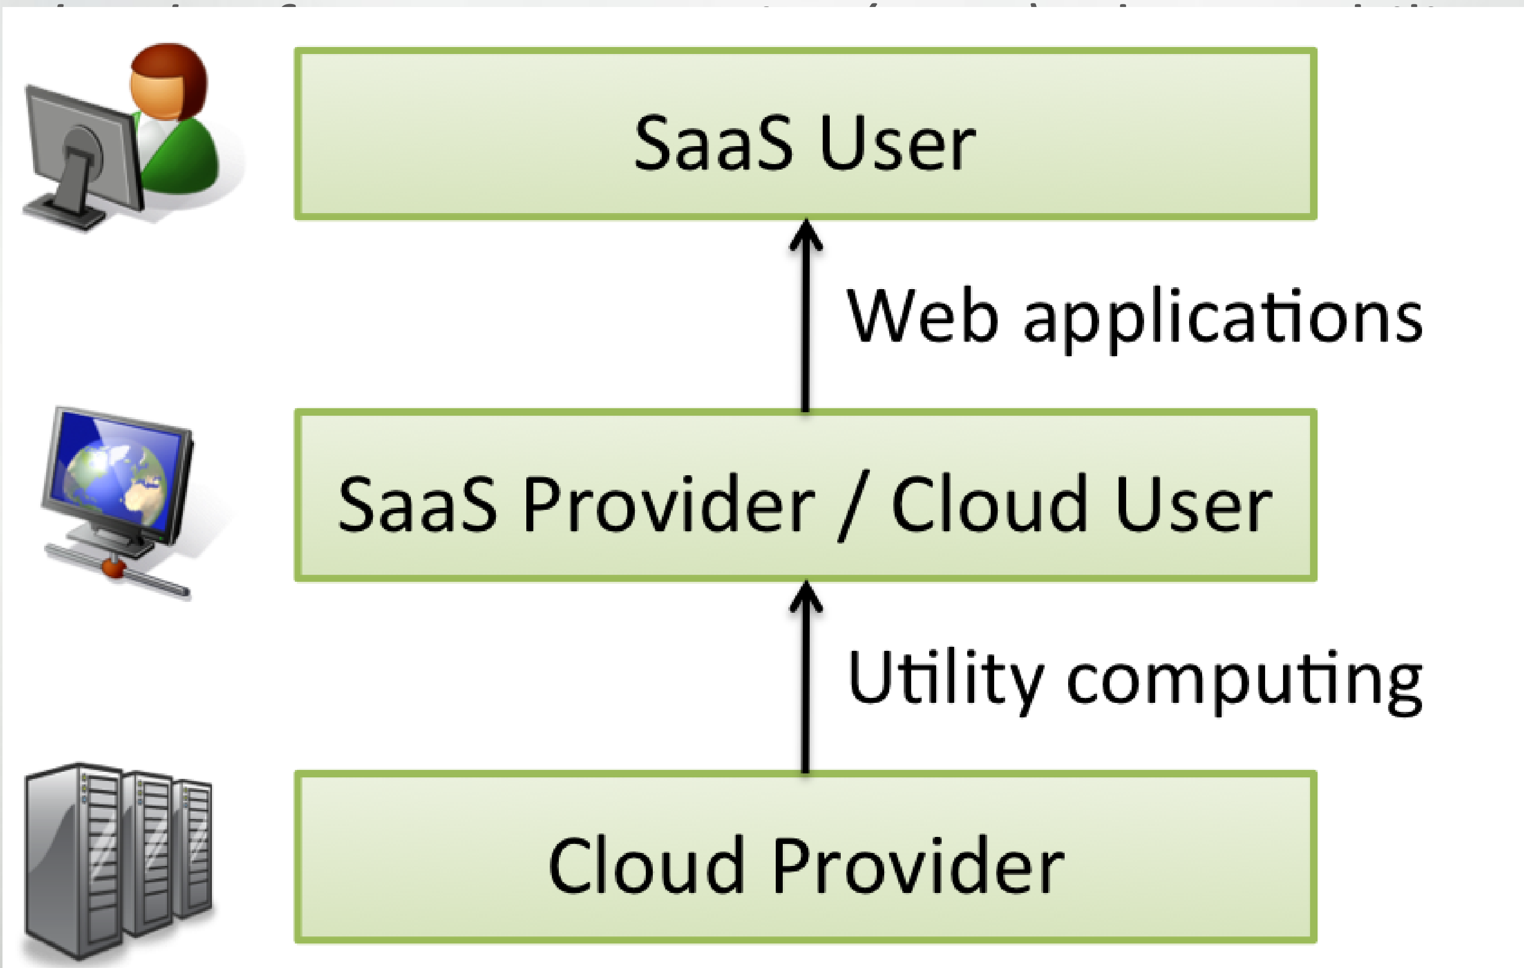

Deployment Models

- *Private cloud.* The cloud infrastructure is owned or leased by a single organization and is operated solely for that organization.
- *Community cloud.* The cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns
- *Public cloud.* The cloud infrastructure is owned by an organization selling cloud services to the general public or to a large industry group.
- *Hybrid cloud*.

What About HPC Moving to the Cloud?

CRAY THE SUPERCOMPUTER COMPANY

"Overall results indicate that EC2 is <u>six</u> times slower than a typical mid-range Linux cluster, and <u>twenty</u> times slower than a modern HPC system. The interconnect on the EC2 cloud platform severely limits performance and causes significant variability."

– United States Department of Energy. *Performance Analysis of High Performance Computing Applications on the Amazon Web Services Cloud.* http://www.nersc.gov/news/reports/technical/CloudCom.pdf

A two generation old Cray supercomputer is more scalable and *20x faster* than Amazon's current cloud—limiting its productivity on today's HPC applications

(This independent study used only 3,500 cores out of over 38,000 on the Cray system)

The Exascale Challenge

Breaking Sustained Performance Barriers

1 GF – 1988: Cray Y-MP; 8 Processors

- Static finite element analysis


1 TF – 1998: Cray T3E; 1,024 Processors

- Modeling of metallic magnet atoms


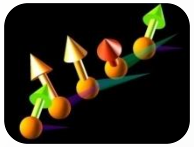

1 PF – 2008: Cray XT5; 150,000 Processors

- Superconductive materials

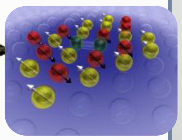

1 EF – ~2018: ~10,000,000 Processors

Technology Trends in Supercomputing

Vector Supercomputing
- Small scale parallelism
- Very fast scalar processing

Symmetric Multiprocessing (SMP)
- Moderate parallelism
- Large shared memory

Massive Parallel Processing
- Large scale parallelism
- Distributed data models
- Networks, Clusters, Grids, Clouds

Manycore/Hybrid
- Extreme parallelism
- Power Efficiency
- Applications Targeted Processing

Expectations of an Exaflop by ~2018-2020

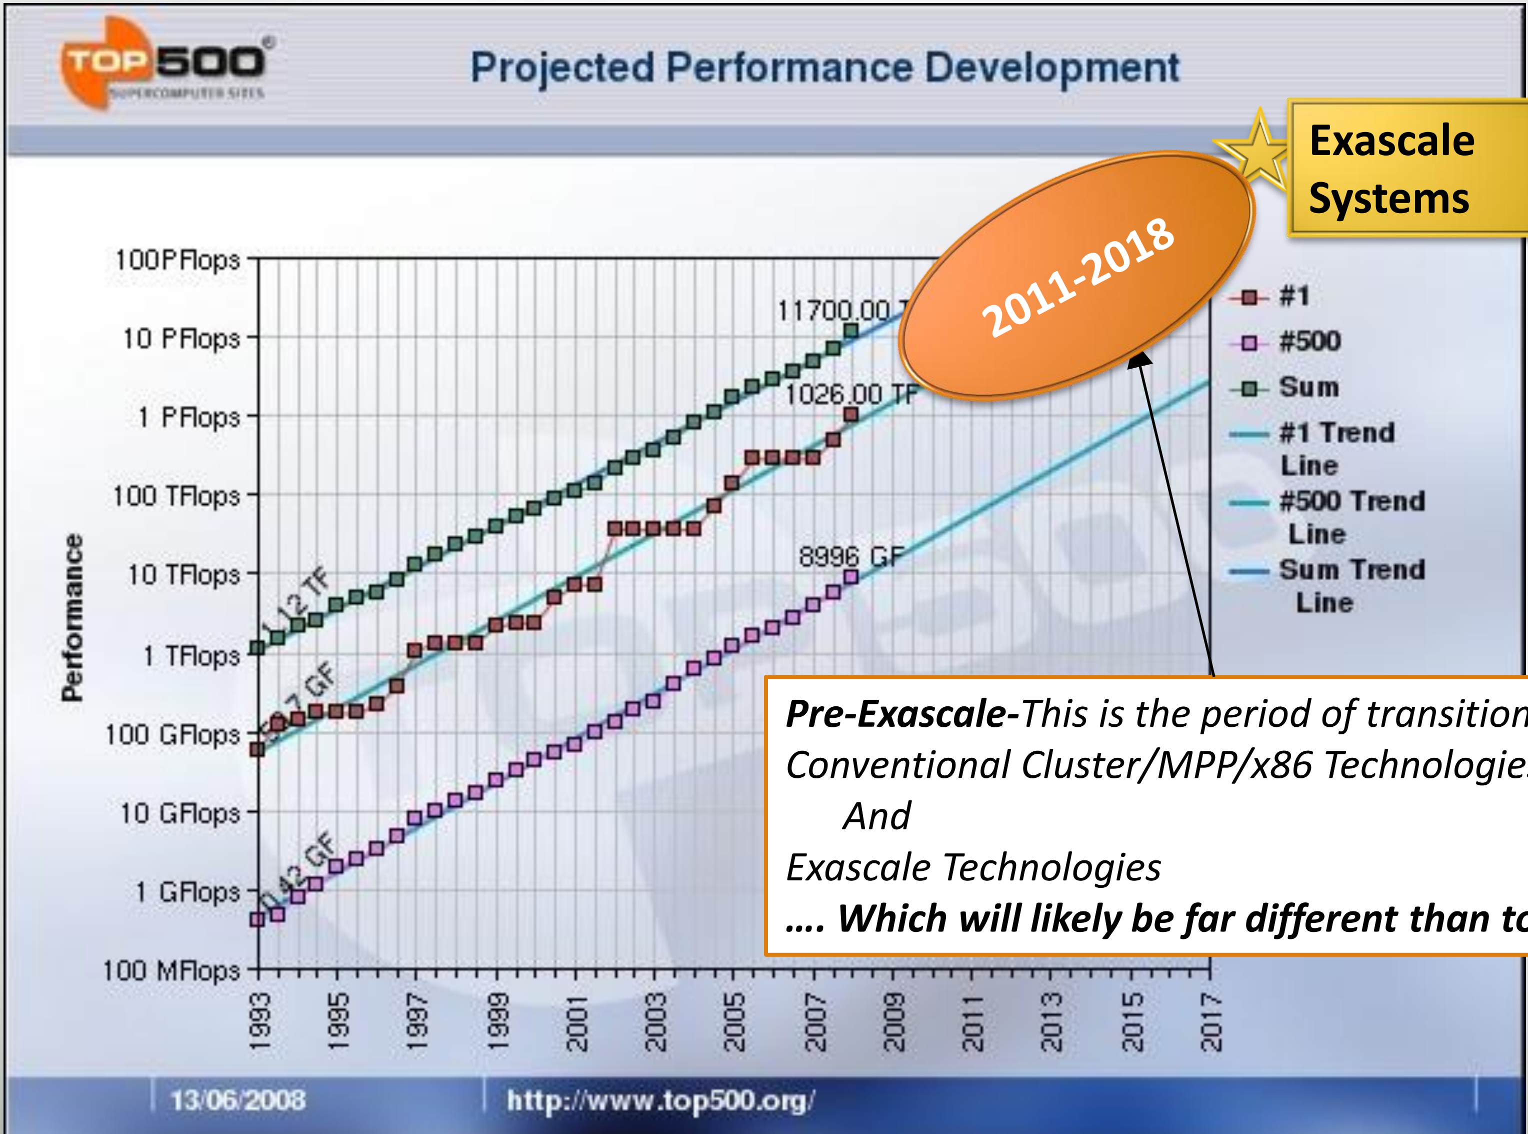

Pre-Exascale-*This is the period of transition between Conventional Cluster/MPP/x86 Technologies*

And

Exascale Technologies

…. Which will likely be far different than today

Increasing Importance of Scaling

- Per-core performance has stalled
- Rate of increase has increased with advent of multi-core chips
- Top systems have more than 100,000 processing cores today
- Million processor systems expected
- Scaling (applications, interconnects, system software, tools, reliability, etc.) is the dominant issue

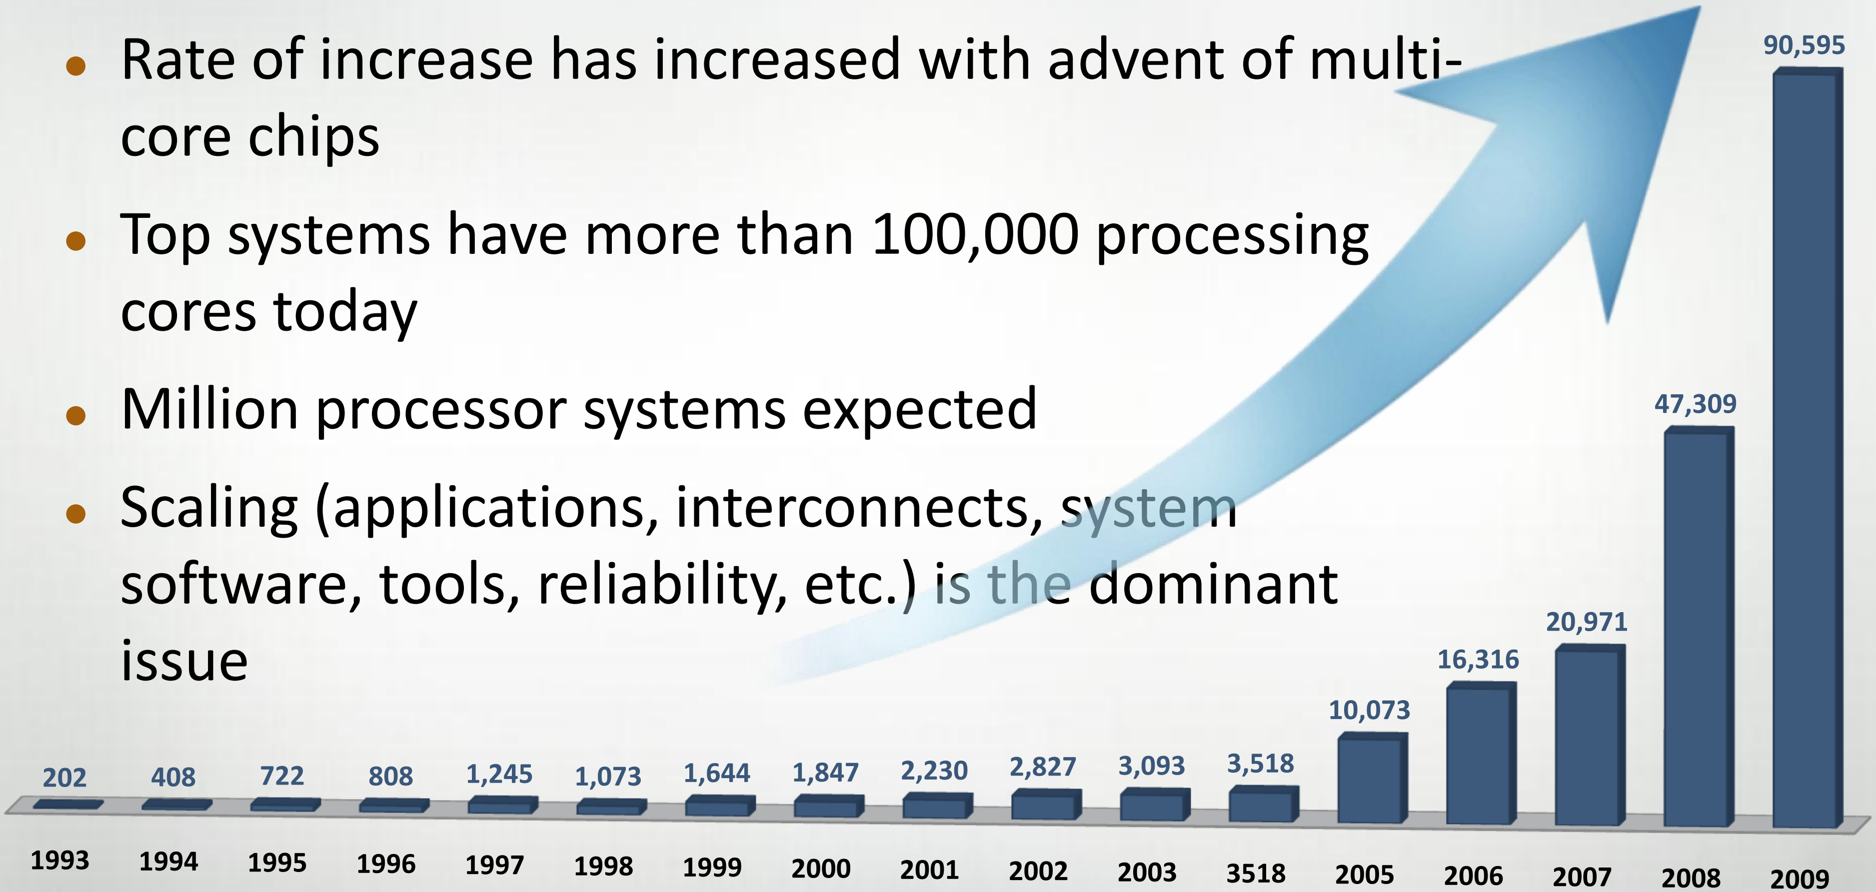

Average Number of Processor Cores per Supercomputer (Top 20 of Top500)

Source: www.top500.org

Exascale Computing: Potential System Architecture

Systems	2009	2019-2020	Difference Today & 2019
System peak	**2 Pflop/s**	**1 Eflop/s**	**O(1000)**
Power	**6 MW**	**~20 MW**	**O(1000)**
System memory	0.3 PB	32 - 64 PB **[.03 Bytes/Flop]**	O(100)
Node performance	125 GF	1,2 or 15TF	O(10) – O(100)
Node memory BW	25 GB/s	2 - 4TB/s **[.002 Bytes/Flop]**	O(100)
Node concurrency	12	O(1k) or 10k	O(100) – O(1000)
Total Node Interconnect BW	3.5 GB/s	200-400GB/s (1:4 or 1:8 from memory BW)	O(100)
System size (nodes)	18,700	O(100,000) or O(1M)	O(10) – O(100)
Total concurrency	225,000	O(billion) [O(10) to O(100) for latency hiding]	O(10,000)
Storage	15 PB	500-1000 PB (>10x system memory is min)	O(10) – O(100)
IO	0.2 TB	60 TB/s (how long to drain the machine)	O(100)
MTTI	days	O(1 day)	- O(10)

Courtesy: Jack Dongarra – PARA 2010 Keynote Presentation

Challenge of Exascale

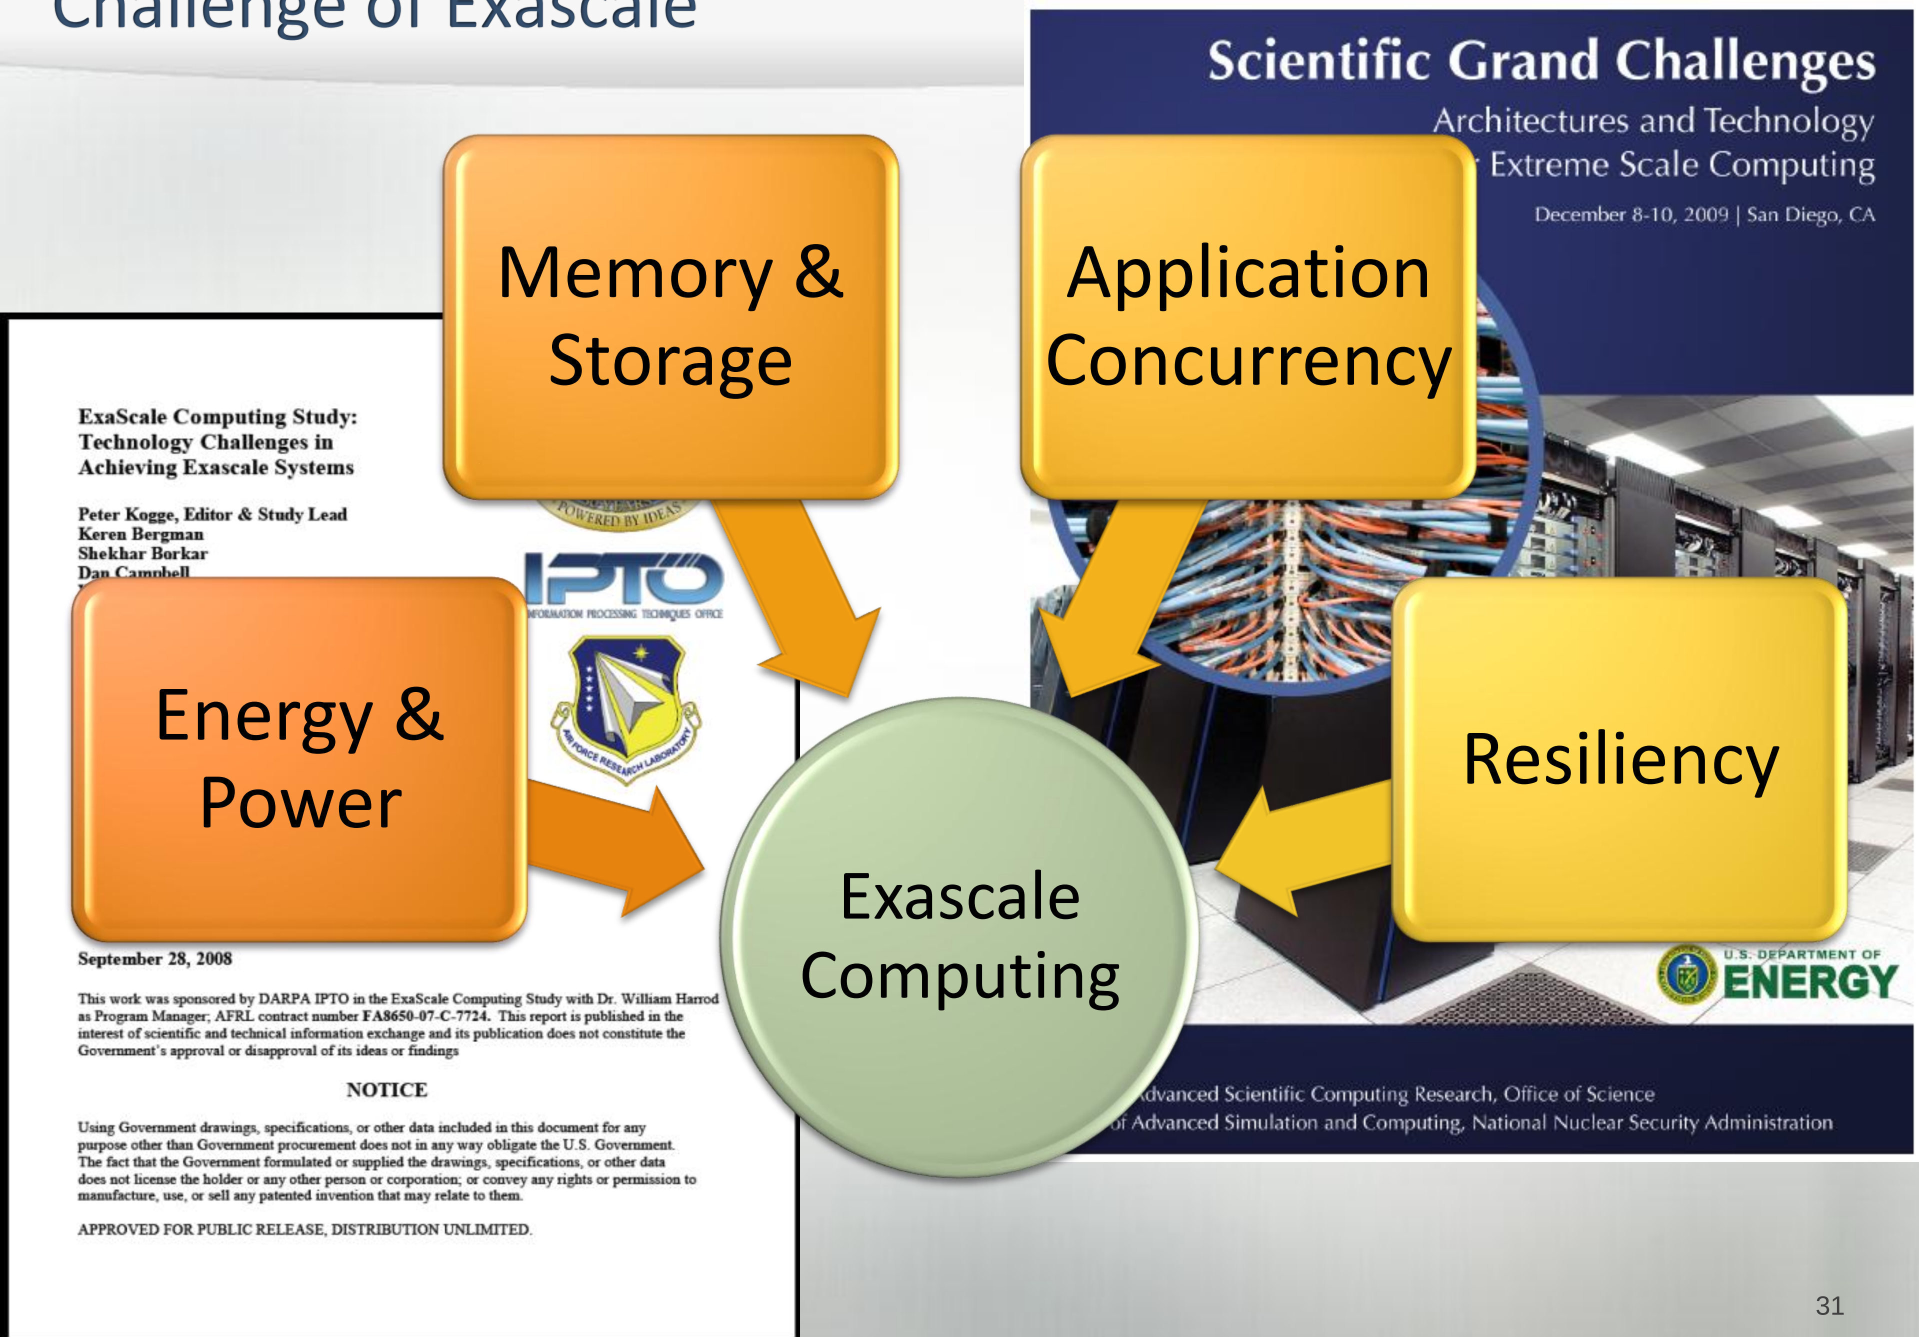

Challenge of Exascale

Memory & Storage
Stacked Memory
On-Chip memory
Storage Appliances
Memory Error Detection and Correction

Application Concurrency
Multithreading
Improved Application Scoping
New concurrent programming models
Network Scalability

Energy & Power
Chip Power States
Lower power memory
Power-Aware Software
Manycore

Resiliency
Software and Hardware to survived component failure
Automatic failover
Network and Node resiliency

Exascale Computing

- DataCenter Innovation
- Strong industry collaboration
- Strong customer collaborations
- Hardware and Software Development and Innovation

Is Linux Sufficient for the Pre-Exascale Era?

LINUX JOURNAL

Cray Releases Highly Scalable, More Inclusive Super-Linux

Apr 20, 2010 By Justin Ryan (/users/justin-ryan) in

We'd all love to have a supercomputer, but sadly, most of us will never have the chance to put that much *umph* in our computing. If you happen to be in the market for a sweet little Linux box with a half-million cores or so, though, Cray may have just what you're looking for.

On Wednesday, Cray Inc. — make... fastest supercomputers in the world — released the third-genera... operating system (http://investors.cray.com/phoenix.zhtml?c=98390&p=i... newsArticle&ID=1412822). What makes it so super special? For one... that's always a good thing.

And Linux is at the root of one of its newest features and bigg... Cray is calling "Cluster Compatibility Mode." Previous versio... Environment (CLE) have not played nicely with third-party... (http://www.hpcwire.com/features/New-Cray-OS-Brings-ISVs-in-for-a-S... considerable cost and effort was required to port programs... proprietary communication protocol.

April 14, 2010

New Cray OS Brings ISVs in for a Soft Landing

by Michael Feldman, HPCwire Editor

HPC wire

...y has never made a big deal about the custom Linux operating system it packages with its XT ...rcomputing line. In general, companies don't like to tout proprietary OS environments since ...nd to lock custom codes in and third-party ISV applicati... the third generation Cray ...vironment (CLE3) that the company ann... ...www.hpcwire.com/offthewire/... ...s-Linux-Operating- ...42934.html) ... ...uting an ISV-friendly

Linux News Today.org

Cray launches new super computer, running Super Linux

Add to del.icio.us Digg this

CRAY RESEARCH, INC.

Apr. 21, 2010

If you're looking for a very fast, super computer with about 520,000 CPU cores in it, Cray Research happens to make just what you might be looking for.

Cray Research, maker of some of the fastest supercomputers in the world, just released the third-generation of its Linux super-operating system. And as you might expect, it's got all the latest and the greatest hardware as well...

First, Super Linux is at the root of one of its newest features and biggest selling points: what Cray is calling "Cluster Compatibility Mode."

Previous versions of the CLE (Cray Linux Environment) have not played nicely with third-party applications. Considerable cost and many efforts

Global Competitiveness

- China (based on IDC data)
  - Outpacing EMEA and US growth rate
    - Grown from 2% to 5% of global HPC in 3-4 years.
  - Two Petascale systems from China are already on the latest Top500
    - No.1 Tianhe
    - No.3 Nebulae
  - Multiple petascale-capable datacenters in place and ready
  - Developing in-country skills in SW and HW development
- EMEA
  - PRACE Initiatives
  - Developing new network technologies

Nebulae's New Home – National Supercomputing Center in Shenzhen, China

Summary

Cray and Technology Collaboration

- Collaborative R&D
  - Partners with companies, universities and national laboratories on R&D activities in all areas
    - PNNL(Multithreading) , SNL(SW & HW), LANL(Networks), LLNL (SW)
    - DARPA (Productive Supercomputing)
    - DOD (SW &HW)
    - Allinea (SW)
    - Many universities and regional computing centers
    - Centers of Excellence
      - Weather and Climate (NOAA, Korea, Brazil)
      - Exascale (Edinburgh)


JAGUAR
APPLIED MECHANICS
Railway Technical Research Institute
HLRS
CRAY XT5m

Thank You!