High Performance Computing Facilities

Barbara Helland
Advanced Scientific Computing Research
DOE Office of Science

U.S. High End Computing Roadmap

- In May 2004, the President's High-End Computing Revitalization Task Force (HECRTF) put forth an HEC plan for United States Federal agencies
  - Production Systems
  - Leadership Systems

- Guided HEC programs and investments at the National Science Foundation's Office of Cyber Infrastructure (OCI) and in Department of Energy's Office of Science (SC)

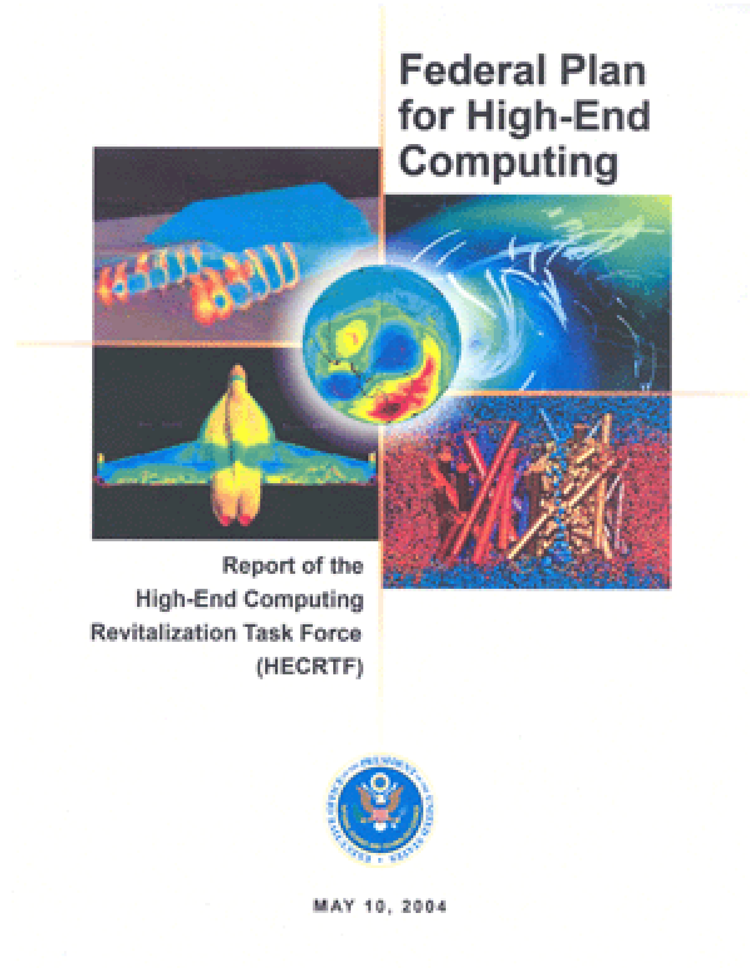

National Science Foundation HPC Strategy

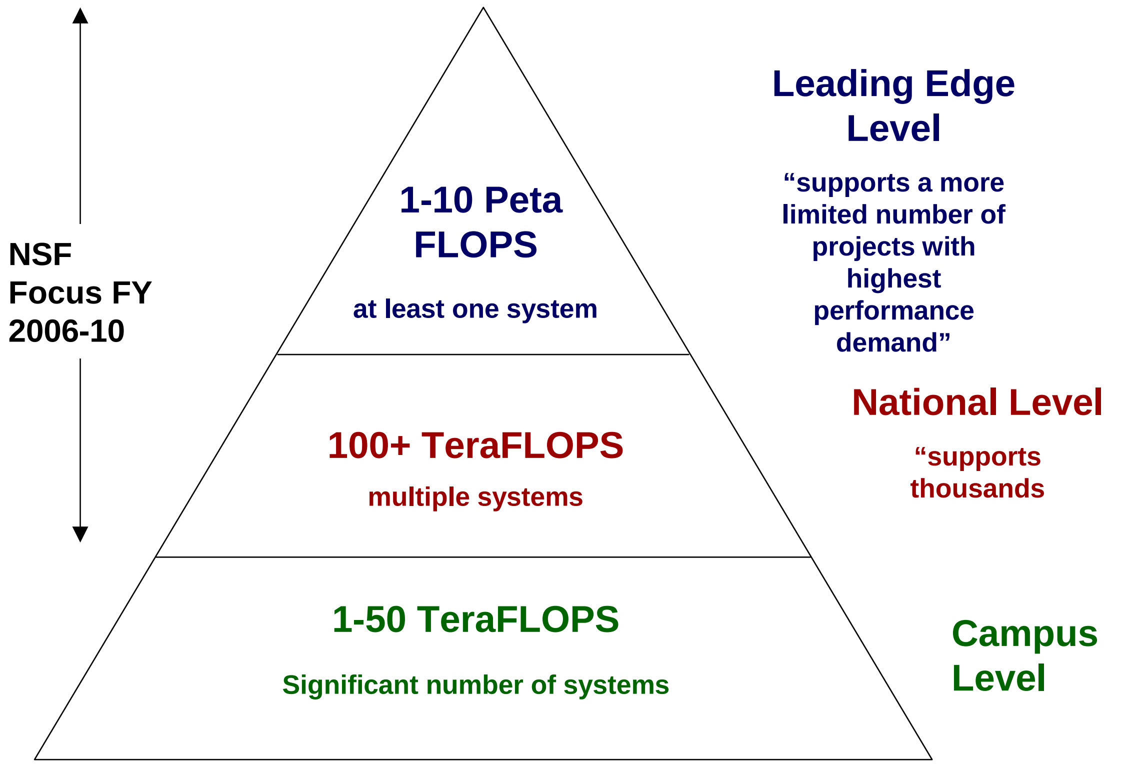

NSF’s goal for high performance computing (HPC) in the period 2006-2011 is to enable petascale science and engineering through the deployment and support of a world-class HPC environment comprising the most capable combination of HPC assets available to the academic community.

NSF National and Leading Edge HPC Resources

National (track-2) resource

- **In Production**
  - **Texas Advanced Computing Center (TACC) – Ranger**
    504 Teraflop (TF) Sun Constellation with 15,744 AMD quad-core processors, 123 terabytes aggregate memory
- **In Process**
  - **University of Tennessee – Kraken**
    - Initial system: 170 TF Cray XT4 with 4,512 AMD quad core processors
    - 1 Petaflop (PF) Cray Baker System to be delivered in 2009

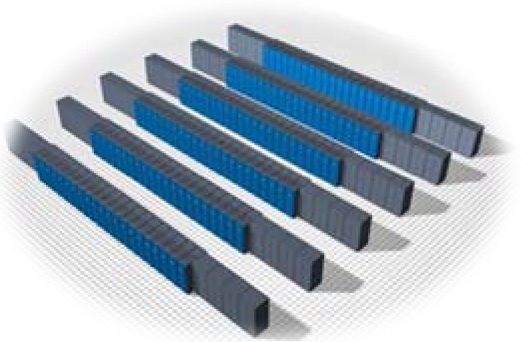

Leading Edge (track-1) resource (in process)

- **National Center for Supercomputing Applications – Blue Waters**
  1 PF sustained performance IBM PERCS with >200,000 cores using multicore POWER7 processors and >32 gigabytes of main memory per SMP

NSF Cyberinfrastructure Allocation Policy

- **Eligibility**
  In general, to apply for an NSF resource allocation you must be a researcher or educator at a U.S. academic or non-profit research institution. A principal investigator (PI) may not be a high school, undergraduate, or graduate student; a qualified advisor must serve in this capacity. A post-doctoral researcher is eligible to serve as PI.
- **Request and Award Sizes**
  - The deadlines and review process for a proposal depends on the amount of resources requested in terms of CPU-hours, or service units (SUs).
  - **Large Resource Allocations Committee (LRAC):** The LRAC reviews proposals requesting more than 200,000 SUs per year across all NSF resources. LRAC meets twice a year, in March and September, and makes awards effective April 1 and October 1, respectively
  - **Medium Resource Allocations Committee (MRAC):** The MRAC reviews proposals requesting between 10,000 and 200,000 SUs per year across all NSF resources. MRAC meets quarterly and makes awards effective January 1, April 1, July 1, and October 1.
  - **Development Allocations Committees (DAC):** Development allocations are requests for up to 10,000 SUs on an individual resource (30,000 SUs on TeraGrid). These allocations target new users planning to submit more substantial proposals in the future, users who have very modest computational requirements, or faculty intending to use resources for classroom instruction or projects.
- **Deadlines**
  - Proposals are due approximately six weeks prior to the meeting at which they will be reviewed.

http://www.ci-partnership.org/Allocations/

NSF Allocations: Petascale Computing

NSF 08-529

- **Program Title:**
  Petascale Computing Resource Allocations (PRAC)

- **Synopsis of Program:**
  In 2011, a new NSF-funded petascale computing system, **Blue Waters**, will go online at the University of Illinois. The goal of this facility is to open up new possibilities in science and engineering by providing computational capability that makes it possible for investigators to tackle much larger and more complex research challenges across a wide spectrum of domains. The purpose of this solicitation is to invite research groups that have a compelling science or engineering challenge that will require petascale computing resources to submit requests for allocations of resources on the Blue Waters system. Proposers must be prepared to demonstrate that they have a science or engineering research problem that requires and can effectively exploit the petascale computing capabilities offered by Blue Waters. Proposals **from or including junior researchers** are encouraged as one of the goals of this solicitation is to build a community capable of using petascale computing.

- Next due date: March 17, 2009

DOE/Office of Science's Advanced Scientific Computing Research Program (ASCR) Facilities Strategy

- High-end Computing – National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory
  - Delivers high-end capacity computing to entire DOE SC research community
- Leadership Computing – Leadership Computing Centers at Argonne National Laboratory and Oak Ridge National Laboratory
  - Delivers highest computational capability to national and international researchers through peer-reviewed **Innovative and Novel Computational Impact on Theory and Computation (INCITE)** program
- Investing in the Future - Research and Evaluation Prototypes
- Linking it all together – Energy Sciences Network (ESnet)

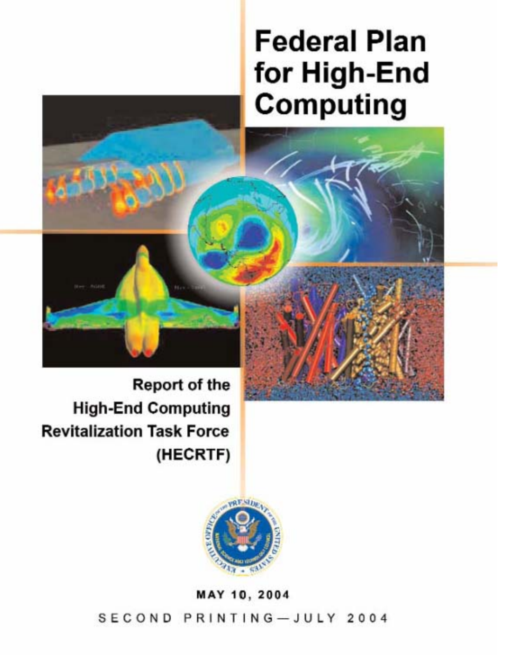

ASCR High Performance and Leadership Computing Facilities

- **NERSC**
  - 104 teraflop Cray XT4 with approximately 9,600 dual core processors; **will upgrade to approximately 360 teraflops with quad corepr processors in Summer, 2008**
  - 6.7 teraflop IBM Power 5 (Bassi) with 888 processors, 3.5 terabytes aggregate memory
  - *3.1 teraflop LinuxNetworx Opteron cluster (Jacquard) with 712 processors, 2.1 terabytes aggregate memory*

- **LCF at Oak Ridge**
  - 263 teraflop Cray XT4 (Jaguar) with 7,832 quad core processors 2.1 GHz AMD Opteron processor nodes, 46 terabytes aggregate memory
  - *18.5 teraflop Cray X1E (Phoenix) with 1,024 multi-streaming vector processors*
  - **Delivery of 1 Petaflop Cray Baker expected in late 2008**

- **Argonne LCF**
  - *5.7 teraflop IBM Blue Gene/L (BGL) with 2,048 PPC processors*
  - 100 teraflop IBM Blue Gene/P began operations April 1, 2008
  - **IBM Blue Gene/P upgraded to 556 TF in March, 2008; upgrade currently in transition to operations**

DOE Allocations

- Facility Reserves
  - 10% of processor time reserved for Director, Office of Science
  - 10% of processor time reserved for Center Director

- NERSC Process
  - 10% of processor time reserved for INCITE
  - 70% of processor time allocated based on the needs of the Program Offices approved by SC-2
    - Large number of projects (350 - 400)
    - Large number of users (2500 – 3100)
    - Single year allocations
    - Call for Allocation Year 2009: August 18 – October 1, 2008

http://www.nersc.gov/nusers/accounts/allocations/

Innovative and Novel Computational Impact on Theory and Experiment- INCITE

- 80% of Leadership Computing resources and 10% of NERSC are allocated through INCITE
- Provides Office of Science computing resources to a small number of computationally intensive research projects of large scale, that can make high-impact scientific advances through the use of a large allocation of computer time and data storage
- Open to national and international researchers, including industry
- No requirement of DOE Office of Science funding
- Peer and computational readiness reviews
- Startup allocation grants available at each site for scaling studies

INCITE 2008

- Received 88 new and 24 renewal proposals requesting over 600 Million processor hours
  - 44% from Universities
  - 46% funded from non-DOE sources
- Over 265 Million hours awarded to 55 new and renewal projects

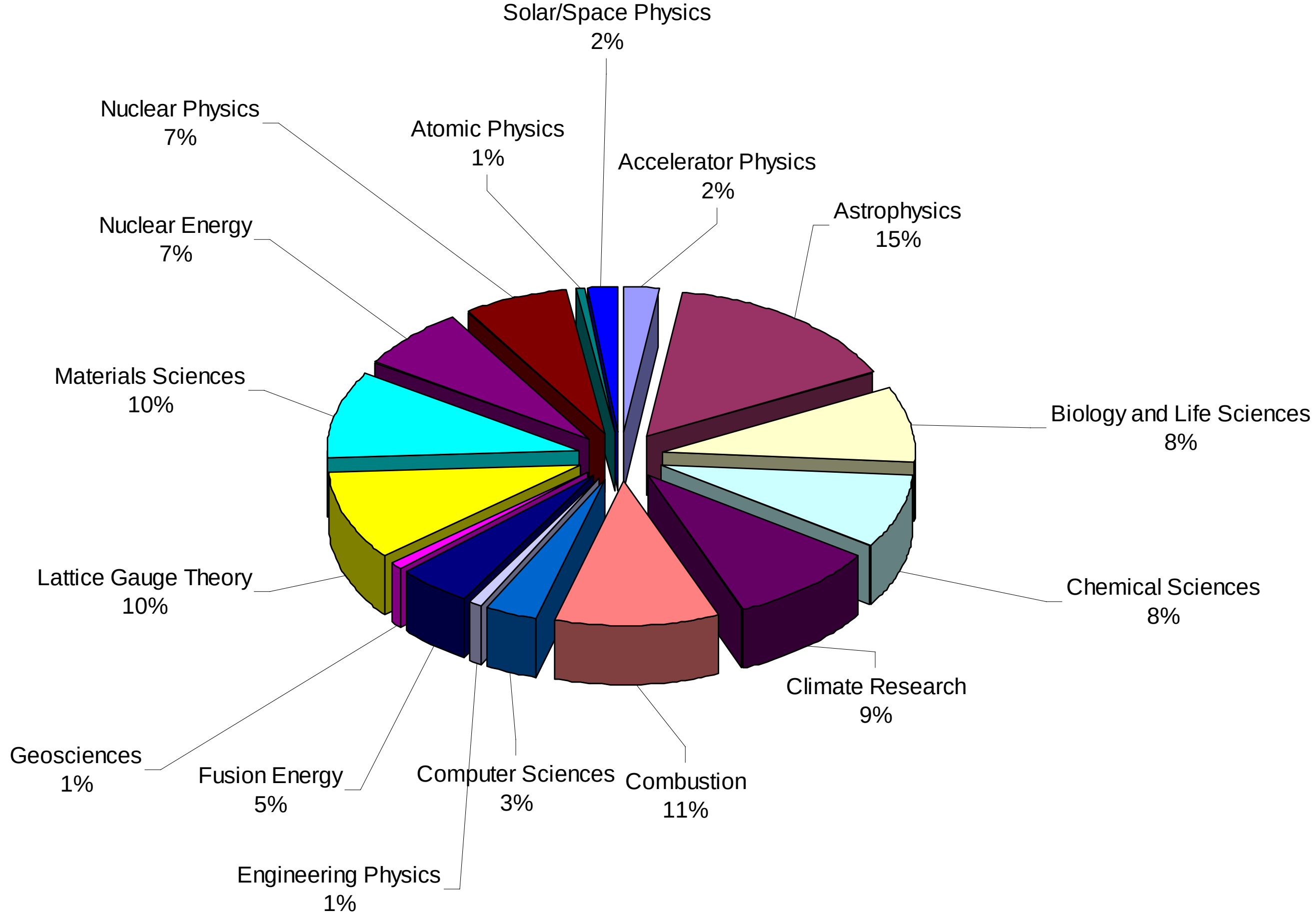

2009 INCITE

New 2009 Call for Proposals for over a half of a billion processor hours of INCITE allocations announced May 13, 2008 at

http://hpc.science.doe.gov

Proposal deadline 11:59 pm EDT August 11, 2008

U.S. Department of Energy
Office of Science

Facilities of the Future

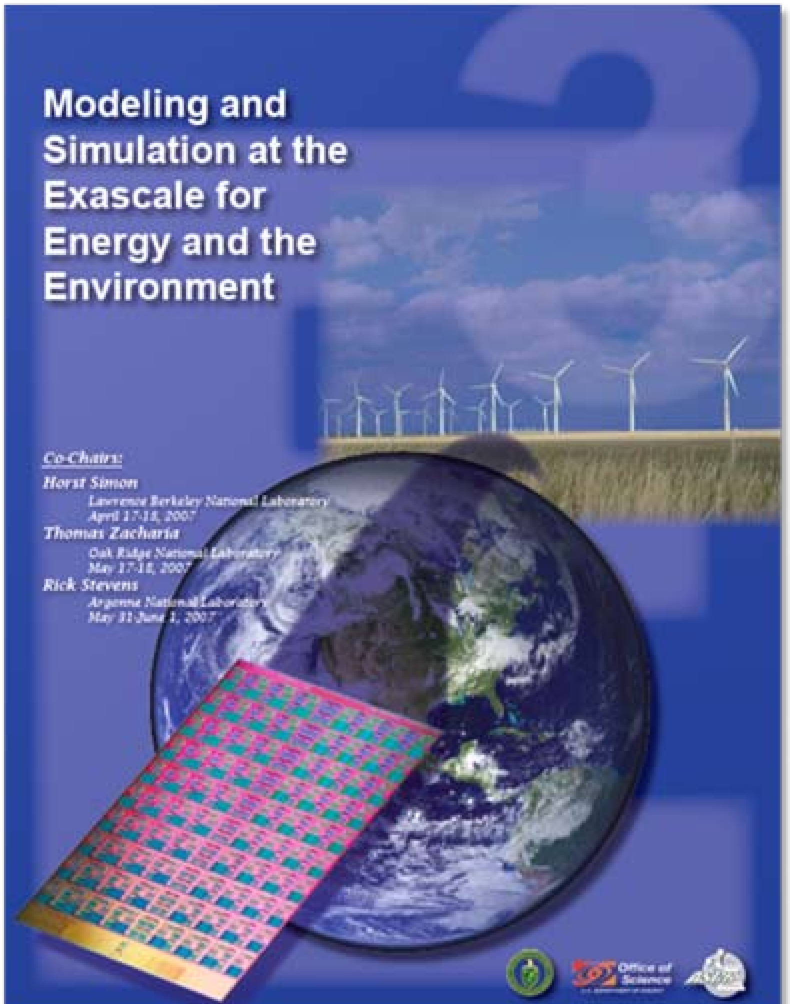

During the spring of 2007 Argonne, Berkeley and Oak Ridge held three Town hall meetings to chart future directions

- Exascale Computing Systems
    - Hardware Technology
    - Software and Algorithms
    - Scientific Applications
- Energy
    - Combustion
    - Fission and Fusion
    - Solar and Biomass
    - Nanoscience and Materials
- Environment
    - Climate Modeling
    - Socio-economics
    - Carbon Cycle

Looking to Exascale

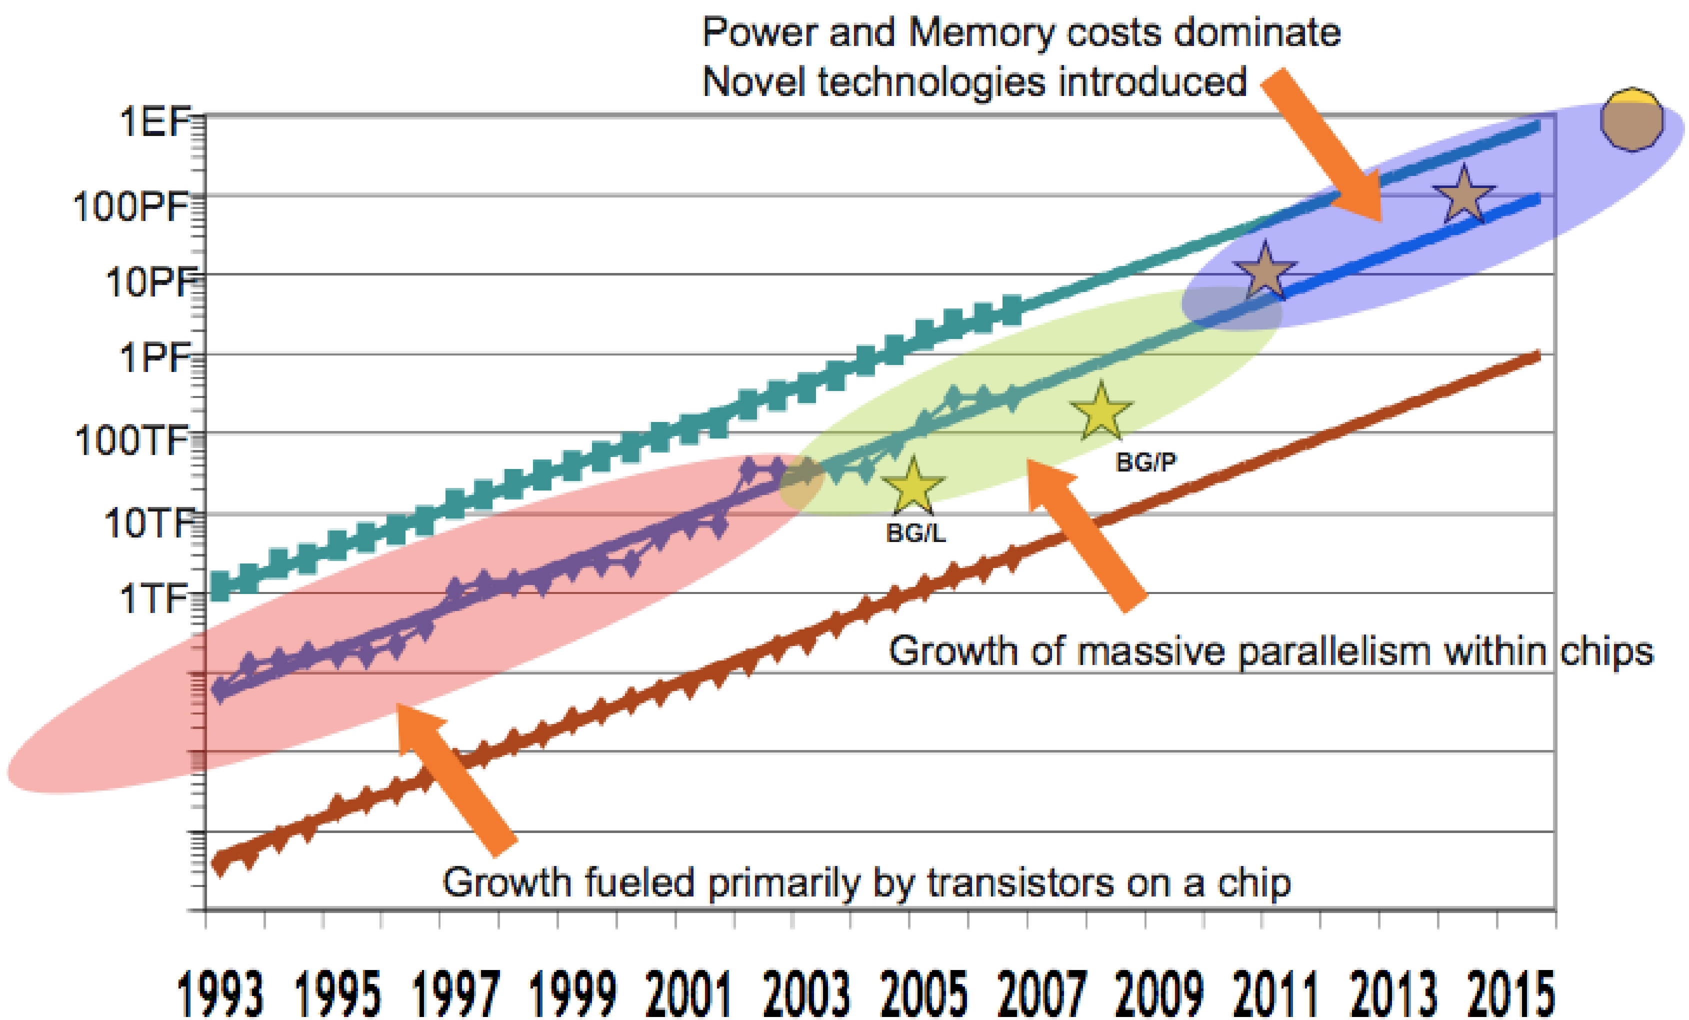

Slide courtesy of Rick Stevens, ANL June 2008

A Three Step Path to Exascale

Begin Full System Delivery (Yr)	2004	2007	2012	2015	2019
Design Parameters	**BG/L**	**BG/P**	**ONE**	**TWO**	**THREE**
Cores / Node	2	4	8-24	32-64	96-128
Clock Speed (GHz)	0.7	0.85	1.6-4.1	2.3-4.8	2.8-6.0
Flops / Clock / Core	4	4	8-32	8-32	16-64
Nodes / Rack	1024	1024	100-512	256-1024	256-1024
Racks / Full System Config	64	72	128-350	128-400	256-400
MB RAM/core	256	512	1024-4096	1024-4096	1024-4096
Total Power	2.5MW	4.8MW	8MW-20MW	20MW-50MW	40MW-80MW
Flops / Node (GF)	5.6	14	128-640	640-2000	2000-6000
Flops / Rack (TF)	5.7	14	200-400	400-1200	1600-4800
LB Concurrency	5.E+05	1.E+06	1M-2M	10M-100M	400M-100M
Full System					
Total Cores (Millions)	0.13	0.3	.3M-1.2M	1M-10M	4M-30M
Total RAM (TB)	33.6	151	2,000-4,400	3,000-10,000	5,000-25,000
Total Racks	64	72	128-350	128-400	256-400
Peak Flops System (PF)	**0.37**	**1**	**25**	**300**	**1200**

Slide courtesy of Rick Stevens, ANL June 2008

Will you be ready?

Backup

DARPA High Productivity Computing Systems (HPCS)

Goal:

- **Provide a new generation of economically viable high productivity computing systems for the national security and industrial user community (2010)**

Impact:

- **Performance** (time-to-solution): speedup critical national security applications by a factor of 10X to 40X
- **Programmability** (idea-to-first-solution): reduce cost and time of developing application solutions
- **Portability** (transparency): insulate research and operational application software from system
- **Robustness** (reliability): apply all known techniques to **protect against outside attacks**, hardware faults, & programming errors

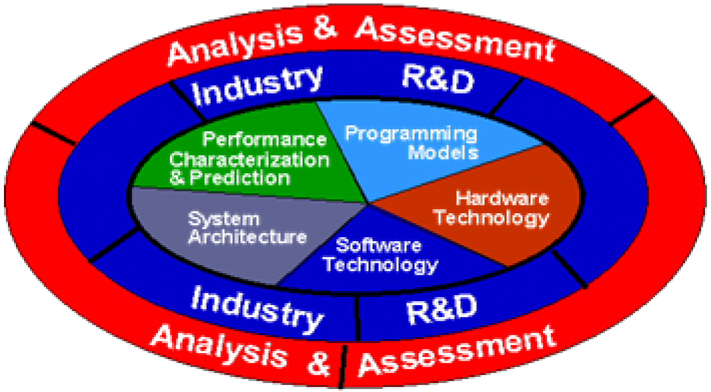

HPCS Program Focus Areas

Applications:

- Intelligence/surveillance, reconnaissance, cryptanalysis, weapons analysis, airborne contaminant modeling and biotechnology

Fill the Critical Technology and Capability Gap

Today (late 80's HPC technology)…..*to*…..Future (Quantum/Bio Computing)

U.S. Department of Energy

Office of Science

HPCS Program Phases I - III

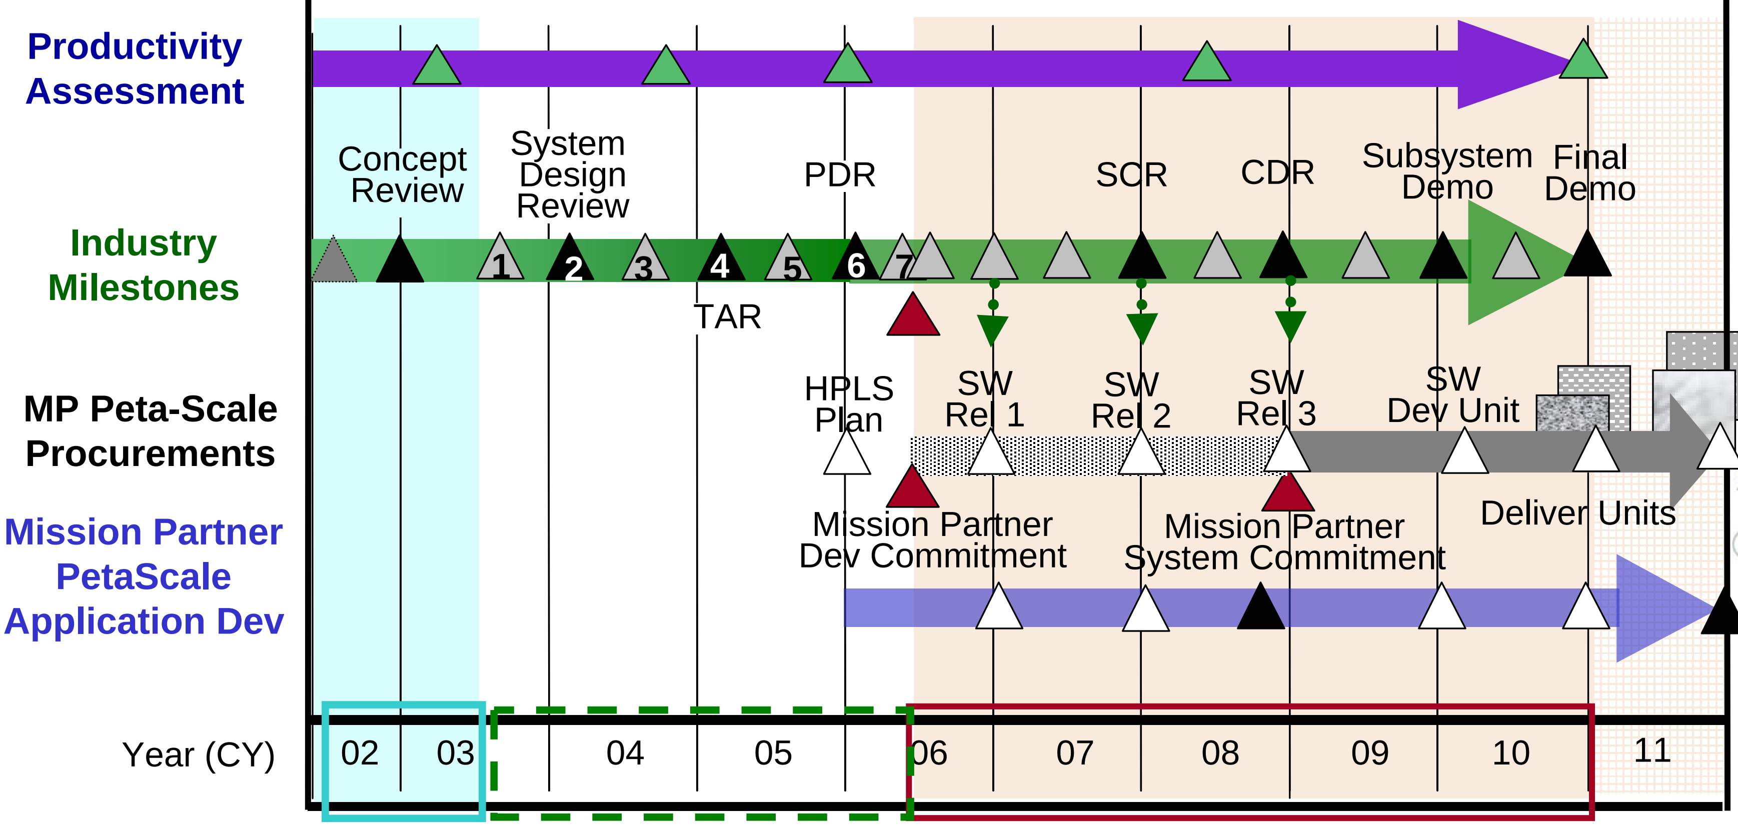

Program Reviews

Critical Milestones

Program Procurements

(Funded Five)
Phase I Industry Concept Study

(Funded Three)
Phase II R&D

Phase III Development and Prototype Demonstration

Mission Partners

CRAY IBM

DOE High-End Computing Revitalization Act of 2004

The Secretary (through the Director of the Office of Science) shall

- Carry out a program of research and development (including development of software and hardware) to advance high-end computing systems
- Develop and deploy high-end computing systems for advanced scientific and engineering applications
- Establish and operate 1 or more Leadership facilities
- Establish at least 1 Software Development Center