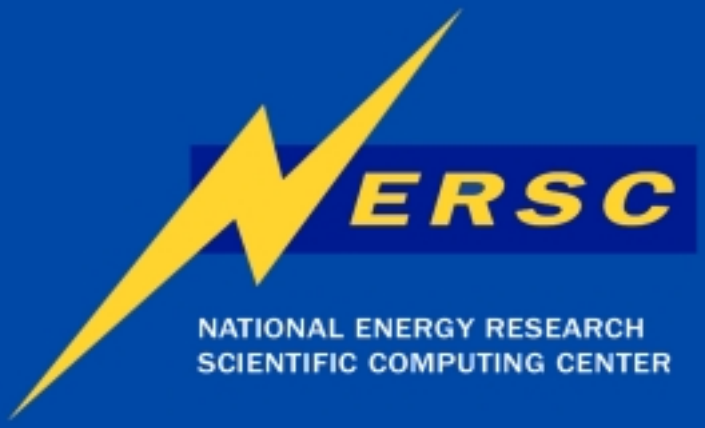

New Computing Initiatives in the US

Horst D. Simon
Director, NERSC Center Division, LBNL
June 25, 2003
http://www.nersc.gov/~simon

Signposts of Change in HPC

In early 2002 there were several signposts, which signal a fundamental change in HPC in the US:

- **Installation and very impressive early performance results of the Earth Simulator System (April 2002)**
- **Lack of progress in computer architecture research evident at Petaflops Workshop (WIMPS, Feb. 2002)**
- **Poor or non-existing benchmarks on sustained systems performance (SSP) for the NERSC workload (March 2002)**

The Earth Simulator in Japan

COMPUTENIK!

- Linpack benchm
  TF/s = 87% of 40
- Completed Apri
- Driven by climat
  earthquake simulation
- Gordon Bell Prize at SC2002

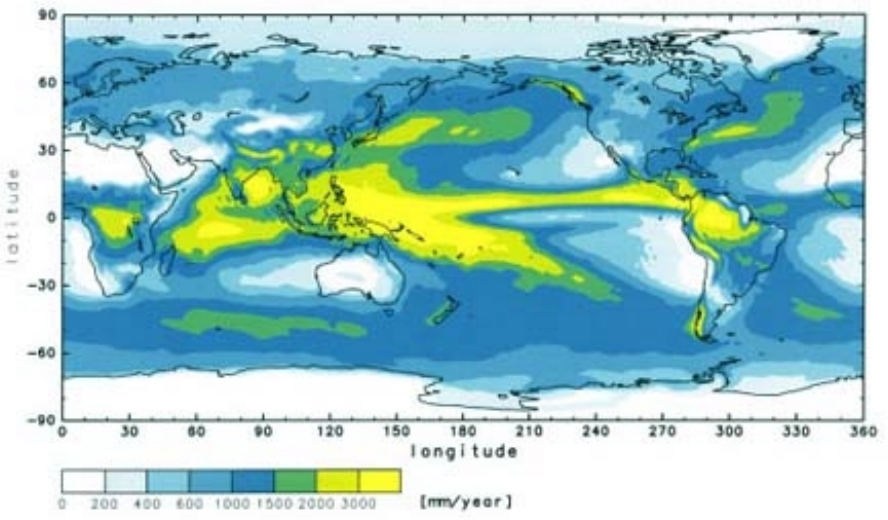

http://www.es.jamstec.go.jp/esrdc/eng/menu.html

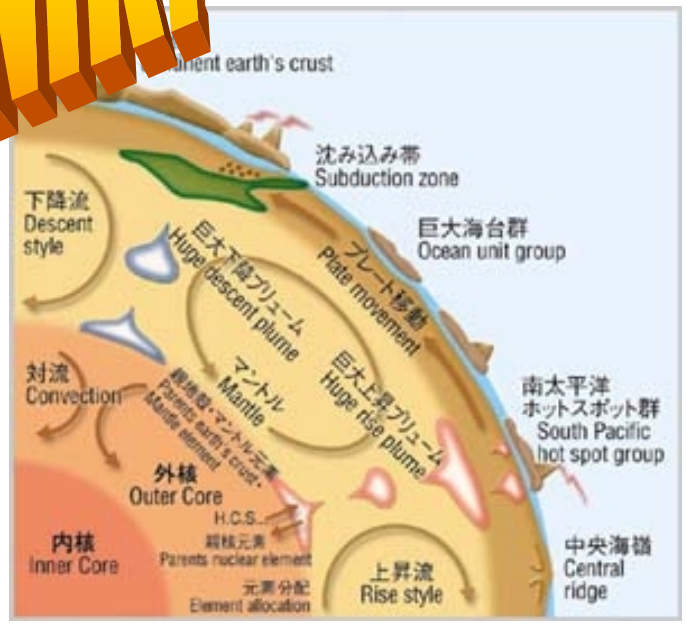

Understanding and Prediction of Global Climate Change	Understanding of Plate Tectonics
Occurrence prediction of meteorological disaster	Understanding of long-range crustal movements
Occurrence prediction of El Niño	Understanding of mechanism of seismicity
Understanding of effect of global warming	Understanding of migration of underground water and materials transfer in strata
Establishment of simulation technology with 1km resolution	

The Divergence Problem

- The requirements of high performance computing for science and engineering and the requirements of the commercial market are diverging.
- The commercial cluster of SMP approach is no longer sufficient to provide the highest level of performance
  - Lack of memory bandwidth
  - High interconnect latency
  - Lack of interconnect bandwidth
  - Lack of high performance parallel I/O
  - High cost of ownership for large scale systems

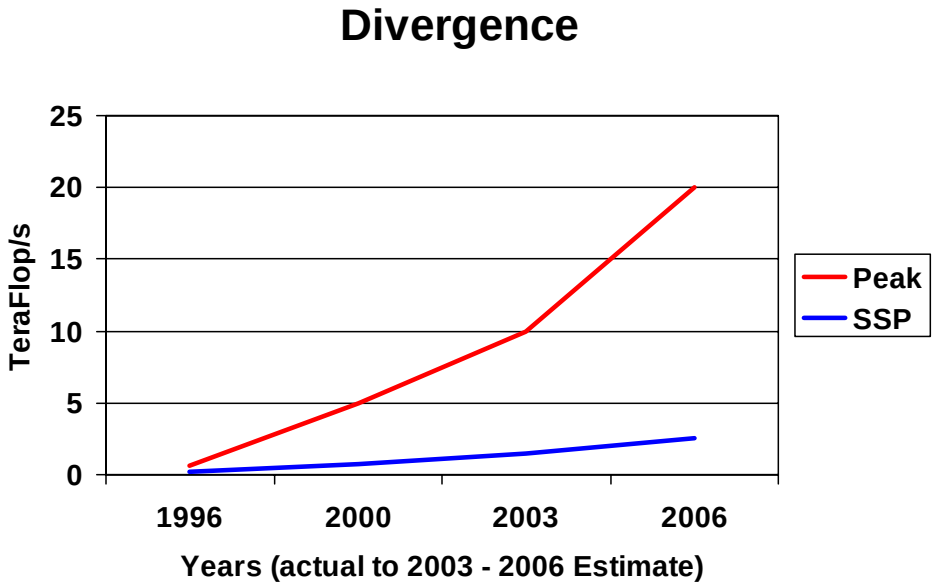

Recent opinions on commodity technology in supercomputing

- "Gordon Bell, now a senior researcher at Microsoft, warns that off-the-shelf supercomputing is a dead end."

  quoted from MIT Technology Review, Feb 2003.

- "Beowulf is dead"

  Thomas Sterling, Caltech, quoted from a panel discussion at SC2002, Nov. 2002

Cray-Sandia Cooperative Development: “Red Storm”

- Collaboration between Sandia Natl. Lab. and Cray (2004)
- True MPP, designed to be a single system
- Distributed memory MIMD parallel supercomputer
- Fully connected 3-D mesh interconnect. Each compute node processor has a bi-directional connection to the primary communication network.
- 108 compute node cabinets and 10,368 compute node processors (AMD Sledgehammer @ 2.0 GHz) ~20 Tflop/s peak
- ~10 TB of DDR memory @ 333 MHz
- 240 TB of disk storage (120 TB per color)
- Less than 2 MW total power and cooling.
- Less than 3,000 square feet of floor space

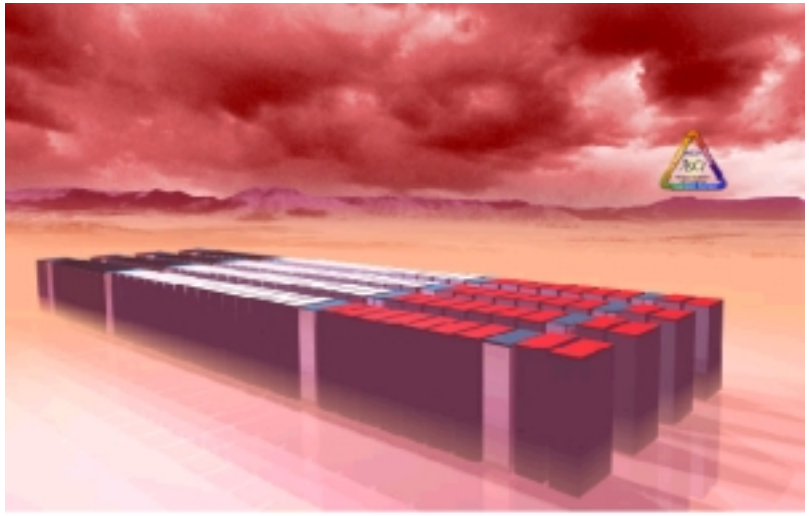

Courtesy: Bill Camp and Jim Thompkins, Sandia

Blue Gene L

(2004: 180 TF system at LLNL)

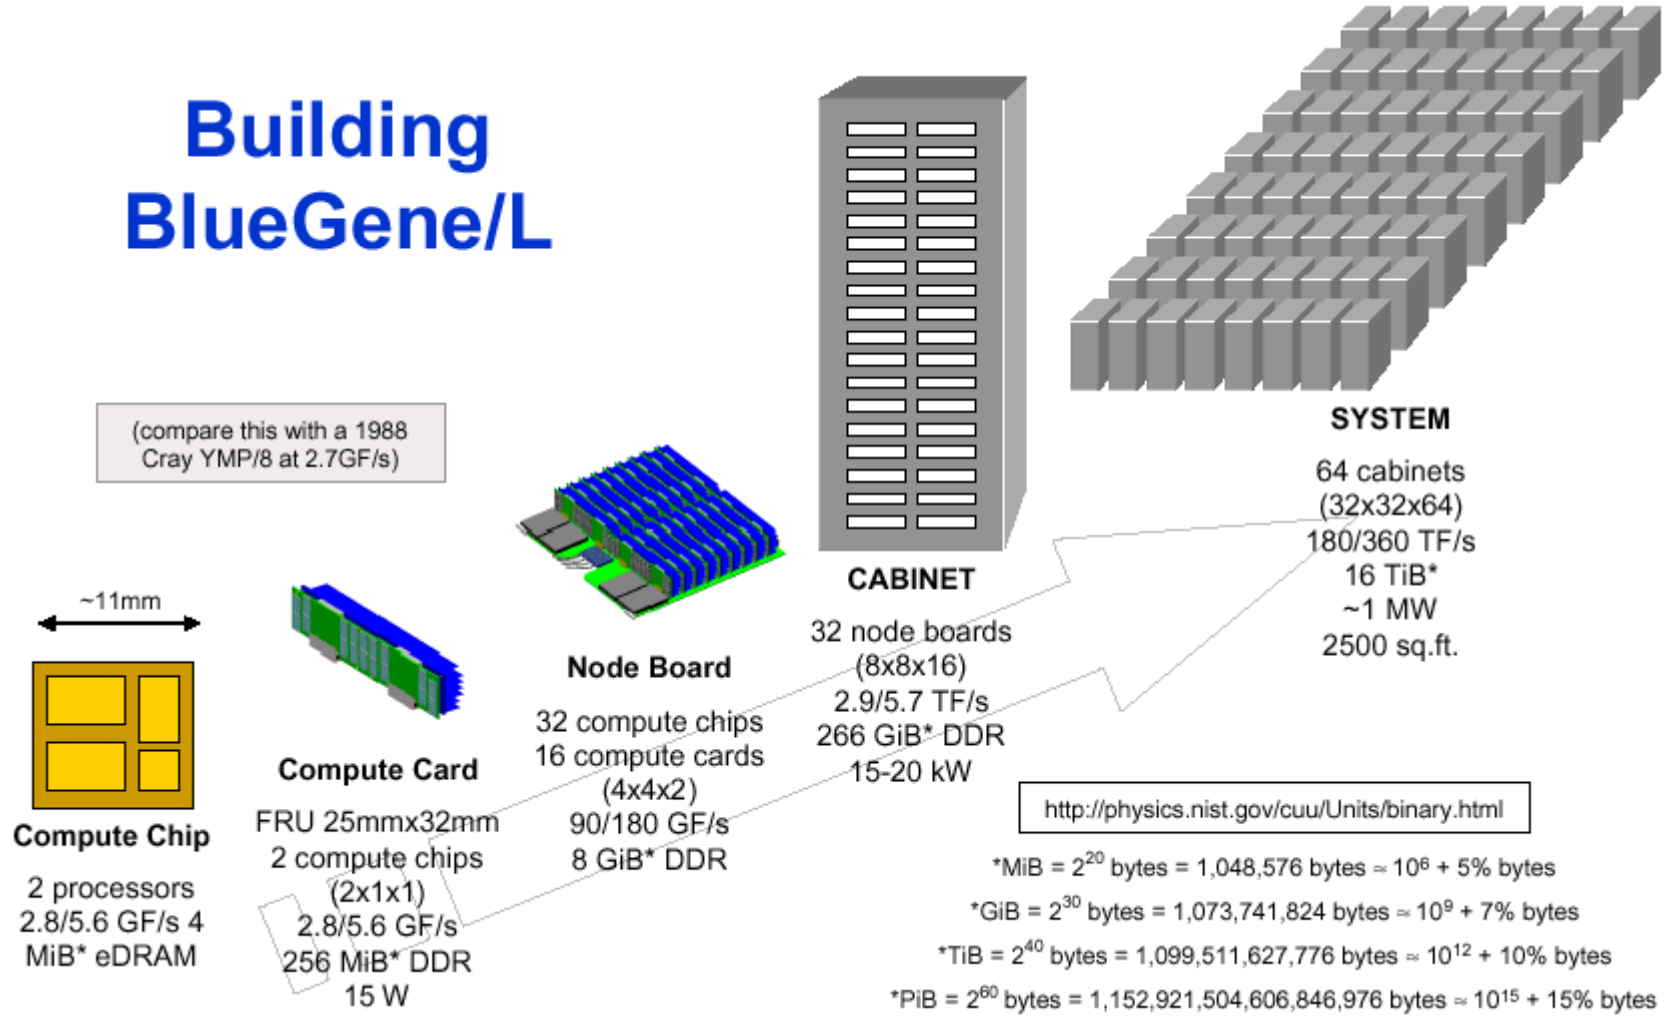

Blue Planet: A Conceptual View
(2005: NERSC and LLNL)

- **Increasing memory bandwidth – single core chips with dedicated caches for 8 way nodes**
- **Increasing switch bandwidth and decreasing latency**
- **Enabling "vector" programming model inside each SMP node**

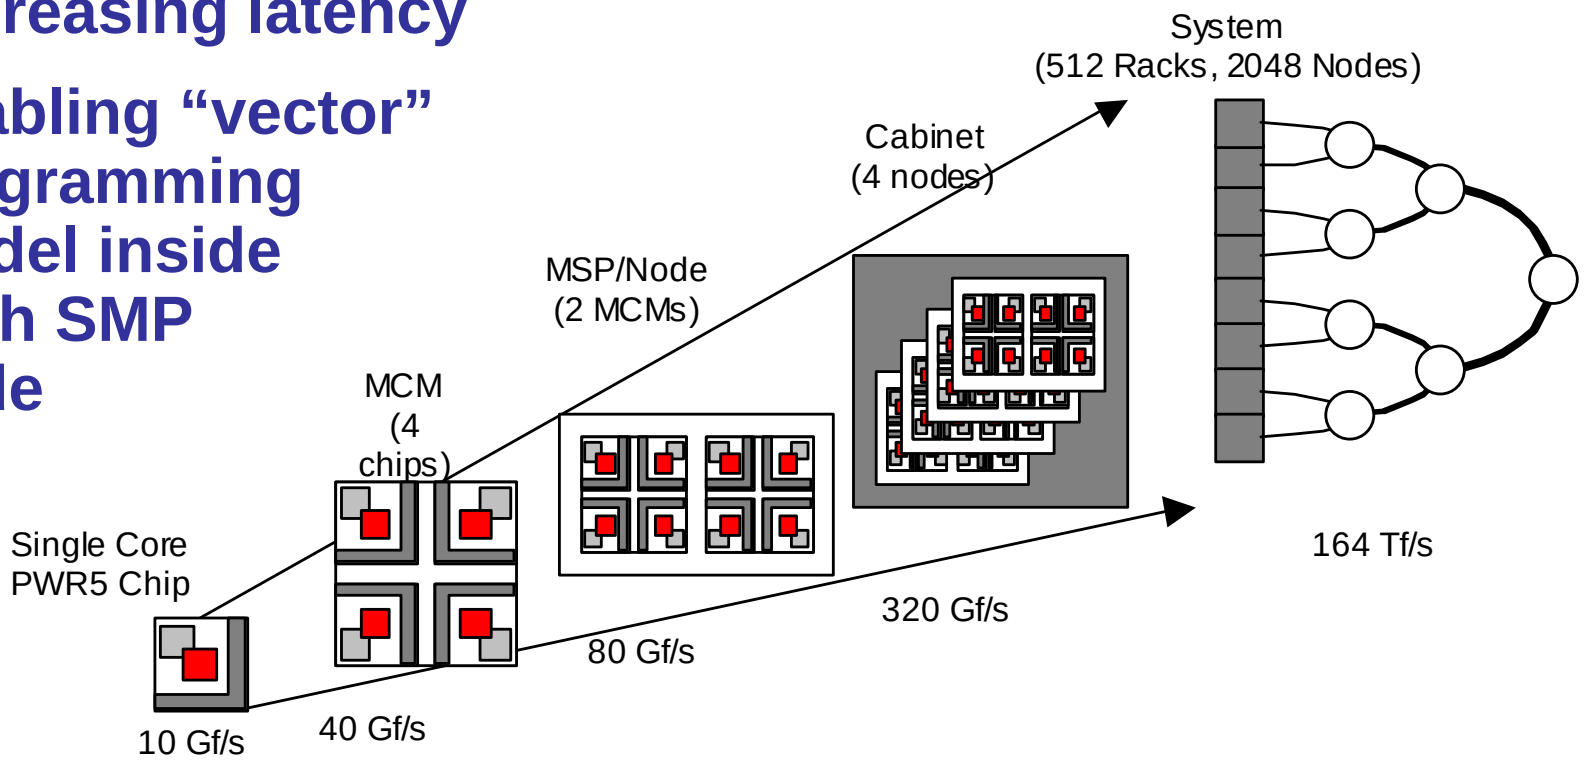

HPCS Planned Program Phases 1-3

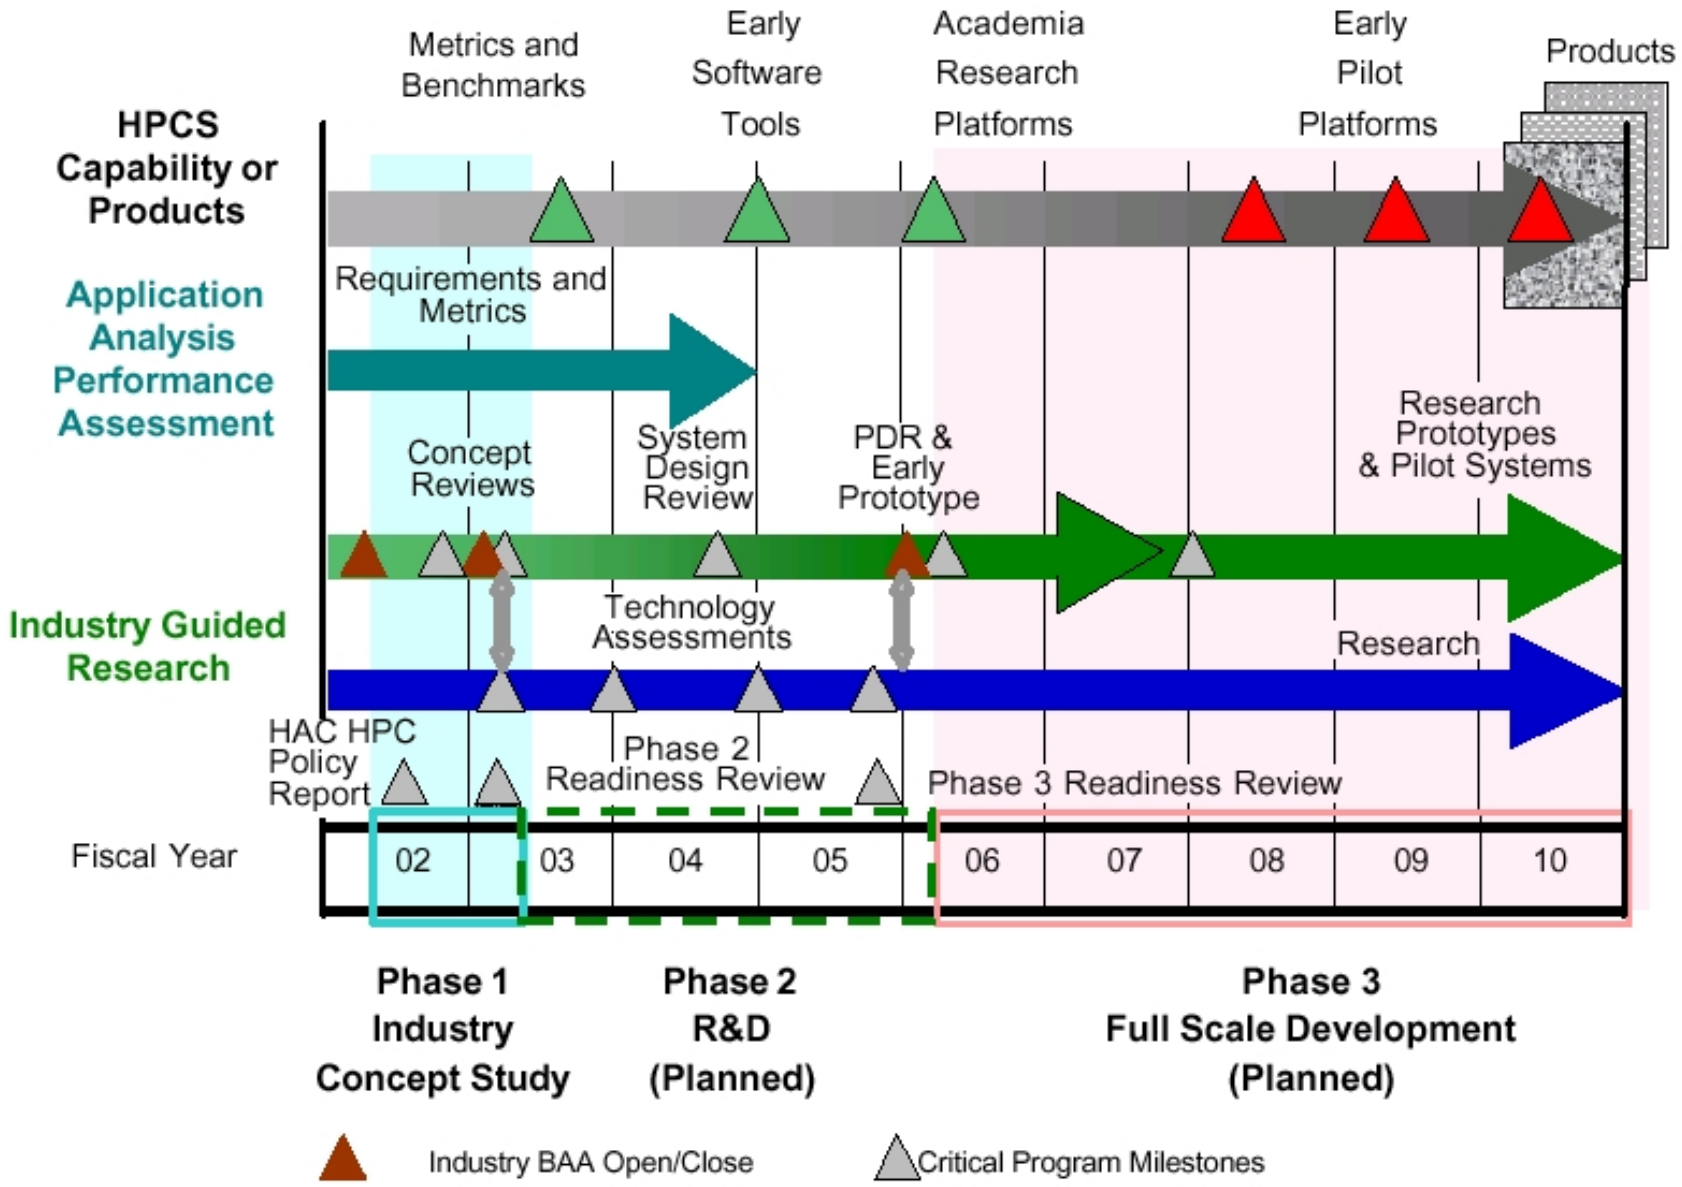

Renewed Interest in High Performance Computing at the Political Level

- **HECRTF (High End Computing Revitalization Task Force)**
  - **OSTP charted; will produce roadmap for all federal agencies**
- **DOE/SC: Ultrascale Initiative (FY2004)**
- **DOE/NNSA: ASCI plans until 2015**
- **NSF: Cyberinfrastructure**
- **DOD: “IHEC” report**
- **NAS: study on “The Future of Supercomputing”**