High-End Computing Interagency Working Group (HEC IWG)

August 16, 2005

John Grosh (DoD)
Chair, HEC IWG

Applications of High-End Computing:
Big Problems with Big Impacts

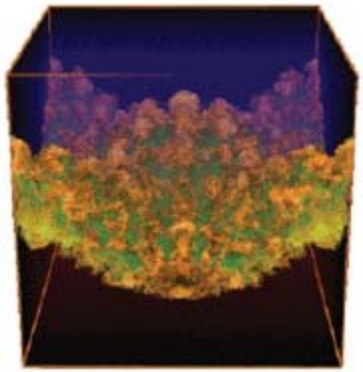

Nuclear Stockpile Stewardship

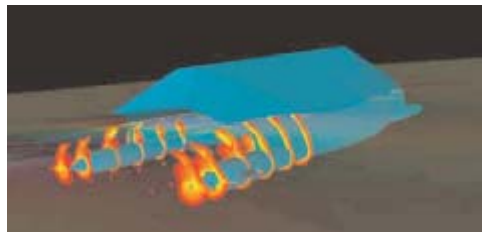

Ship Design

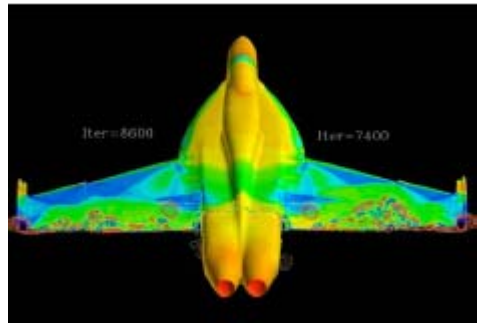

Aeronautics

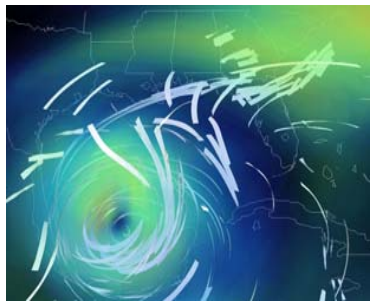

Weather Prediction

Cryptanalysis

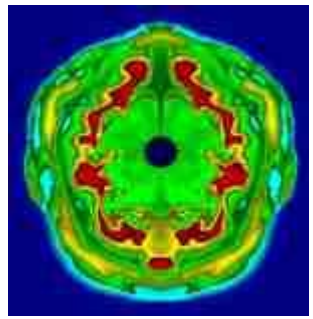

Astrophysical Simulation

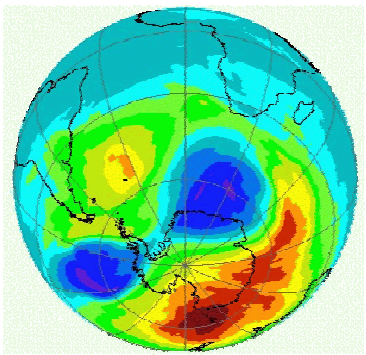

Climate Modeling

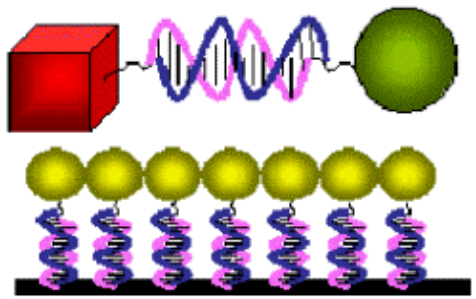

Nano-Science

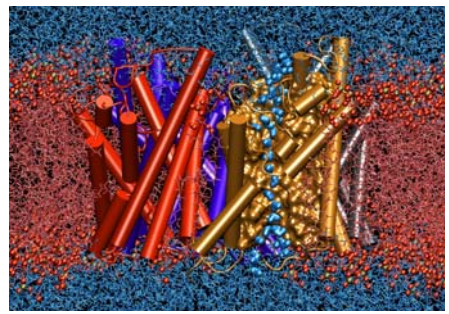

Biology

NITRD Program Coordination

White House
Executive Office
of the President

Office of Science and Technology Policy

U.S. Congress

NITRD Agency
Authorization and Appropriations
Legislation

National Science and Technology Council

Committee on
Technology

Committee on
Homeland and
National Security

National Coordination Office (NCO)
for Networking and
Information Technology
Research and Development

Subcommittee on
Networking and Information
Technology Research and
Development (NITRD)

Subcommittee on
Infrastructure

High End Computing (HEC)
Interagency Working Group

Cyber Security and
Information Assurance (CSIA)
Interagency Working Group

Large Scale
Networking (LSN)
Coordinating Group

High Confidence
Software and
Systems (HCSS)
Coordinating Group

Human Computer
Interaction and
Information
Management (HCI&IM)
Coordinating Group

Software Design and
Productivity (SDP)
Coordinating Group

Social, Economic,
and Workforce
Implications of IT
and IT Workforce
Development (SEW)
Coordinating Group

Agency NITRD Budgets by PCA

FY 2006 Budget Requests (dollars in millions)

Agency	HEC I&A	HEC R&D	HCI&IM	LSN	HCSS	SDP	SEW	Totals
NSF	201.8	105.0	168.5	94.5	76.0	65.4	92.1	803
NIH	135.1	67.3	171.0	76.6	12.3	26.4	12.0	501
DOE/SC	105.7	82.0		36.2			3.5	227
DARPA		81.0	74.4	20.8				176
NSA		36.9		1.5	62.2			101
NASA	34.0		14.5	13.0	12.8			74
AHRQ			38.0	30.0				68
NIST	5.4	0.6	8.6	4.6	18.0	4.8		42
OSD					2.5	20.0		22
NOAA	13.7	1.8	0.5	2.8				20
EPA	3.3		3.0					6
Subtotals	**499.0**	**374.6**	**478.5**	**280.0**	**183.8**	**118.1**	**107.6**	**2041**
DOE /NNSA	33.1	30.5		14.3		31.6	4.4	114
Totals	**532.1**	**405.1**	**478.5**	**294.3**	**183.8**	**149.7**	**112.0**	**2155**

Other participating agencies: AFRL, ARO, DHS, FAA, FDA, NARA, ONR

Federal Plan for High-End Computing

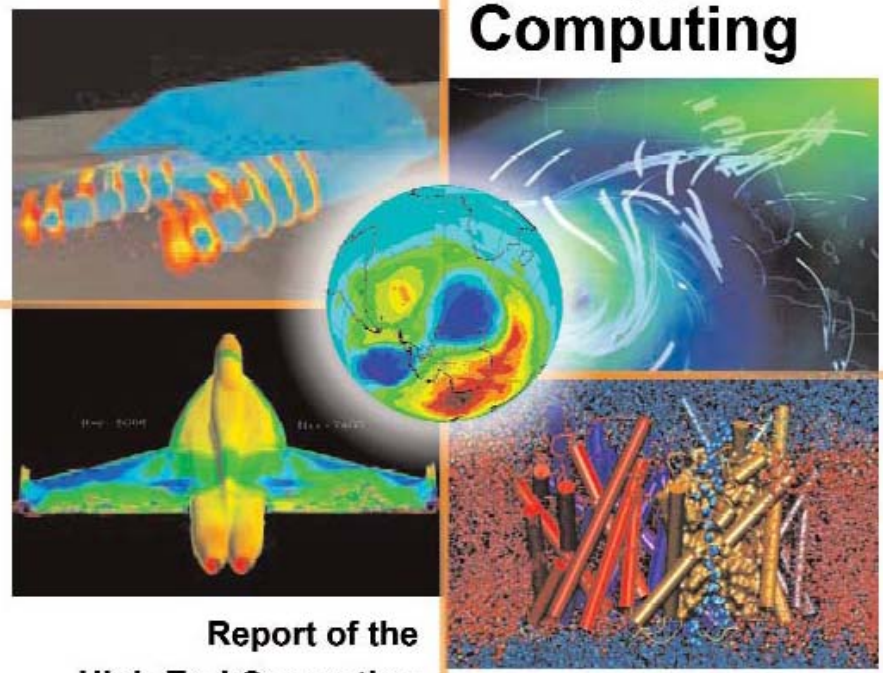

Report of the High-End Computing Revitalization Task Force (HECRTF)

MAY 10, 2004

SECOND PRINTING—JULY 2004

RD&E - Key Technologies

Hardware

- ⇨ Microarchitecture
- ⇨ Memory
- ⇨ Interconnect
- ⇨ Power, cooling, and packaging
- ⇨ I/O and storage

Software

- ⇨ Operating systems
- ⇨ Languages, compilers, and libraries
- ⇨ Software tools and development environments
- ⇨ Algorithms

Systems

- ⇨ System architecture
- ⇨ Reliability, availability, and serviceability (RAS)
- ⇨ System modeling and performance analysis
- ⇨ Programming models
- ⇨ System modeling and performance analysis

Interagency Collaborations in High-End Computing

- **High-End Computing University Research Activity**
  - **FY 2004 Research Areas (22 awards)**
    - Runtime operating systems, languages, compilers, libraries
    - Projects executed out of DOE/SC and NSF
    - Funding provided by DOE/SC, NSF, DARPA, NSA
  - **Agencies plan to continue this effort in FY2006**
- **Leadership Systems (DOE/OS, NASA)**
- **Performance/Productivity Measurement (DARPA, DOE/OS, DOE NNSA, NASA, NSA, NSF, and others)**
- **DARPA High Productivity Computing Systems Program (DOE/OS, DOE NNSA, NSA, NASA, NSF, and others)**

DOE NNSA and Office of Science Recent and Current I/O R&D Focus

- Cluster File Systems that scale for aligned data operations and Symmetric NFS – Lustre, Panasas, PVFS2 (NNSA, Office of Science, OSC, and friends)
- User space file systems (ANL, NNSA)
- Beginning engineering and early R&D for other scalability dimensions
  - Metadata (Lustre, Panasas, PVFS2, NNSA, ANL, UCSC, SDSC)
  - Security (Lustre, Panasas, NNSA, UCSC, U of Minn, CMU, DOD/NSA)
- NFSv4 security, WAN, Performance (NNSA, U of Mich, and industry, NSF)
- pNFS (NNSA, U of Mich, CMU, and industry)
- Middleware to address usage, unaligned I/O, and efficient connection to High Level I/O libraries (ANL, Northwestern, NNSA)
- High Level I/O libraries (ANL, U of Chicago, NCSA, NNSA)
- Early utilization of intelligence near the disk drive (PNL, U of Minn/NNSA)
- Relaxation of Posix Semantics becoming more common (NNSA, ANL, PNL)
- Enterprise Sharing of Global Parallel File System – access from many clusters, (NNSA, UCSC, U of Minn)
- Autonomics (UCSC, NNSA, ANL, CMU)
- New storage devices for hierarchy – mems, mram (UCSC, NNSA, NSF, INSIC)
- Archive (UCSC, U of Minn, NNSA)
- Standards work - POSIX, NFSv4, pNFS, OBSD (NNSA, Office of Science, Industry)

Mostly Plumbing the Stack for Data Parallelism and Nudging Industry Standards for Leverage

DARPA High Productivity Computing Systems (HPCS) Program

Goal:

- **Provide a new generation of economically viable high productivity computing systems for the national security and industrial user community (2009 – 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

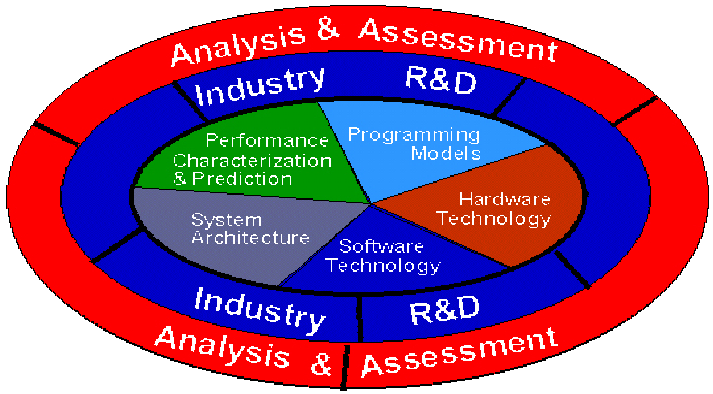

HPCS Program Focus Areas

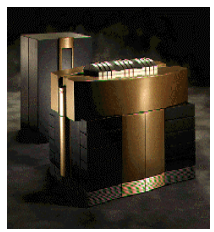

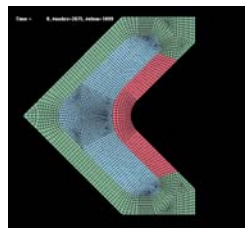

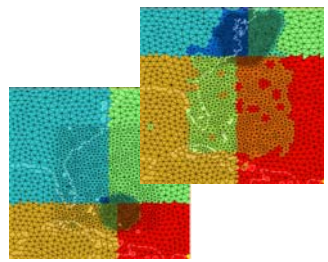

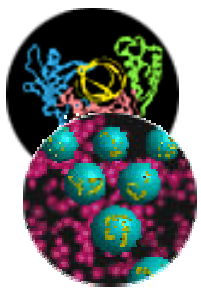

Applications:

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

DARPA HPCS System Architectures

Cray / Sun / IBM

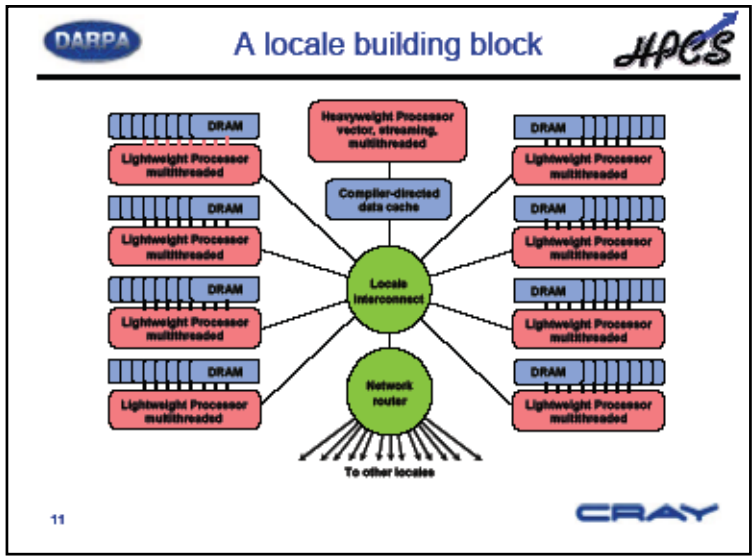

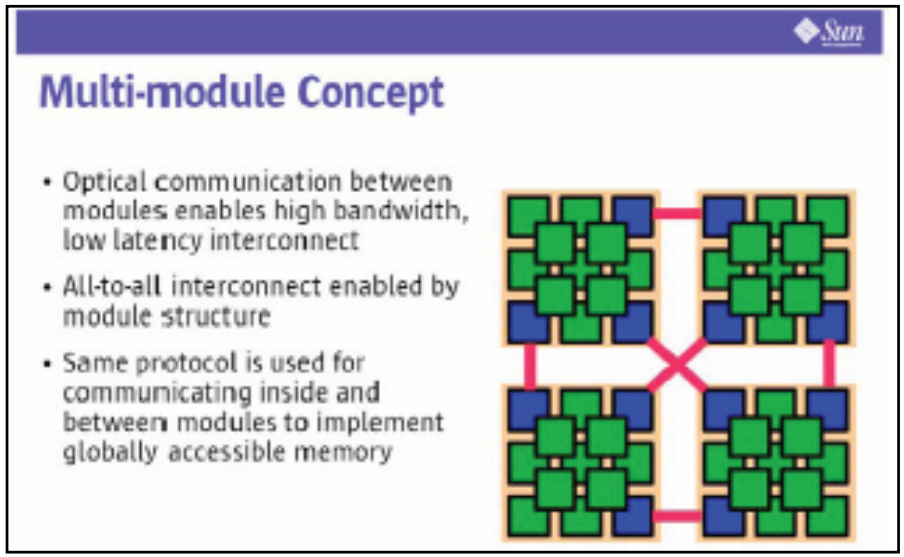

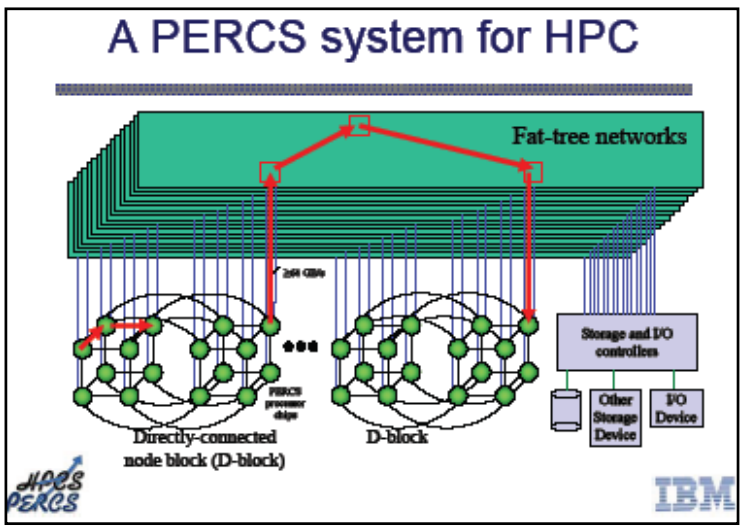

R&D in New Languages

Chapel, X10, Fortress

HPCS I/O Challenges

- 1 Trillion files in a single file system
  - 32K file creates per second
- 10K metadata operations per second
  - Needed for Checkpoint/Restart files
- Streaming I/O at 30 GB/sec full duplex
  - Needed for data capture
- Support for 30K nodes
  - Future file system need low latency communication

An envelope on HPCS Mission Partner requirements

Observations

- **I/O and storage has emerged (? or !)**
- **Workshop represents an opportunity**
  - DARPA HPCS is an example of a program that was focused to address critical National Security need
- **Acknowledgements**
  - **NNSA/ASC Program**
  - **Gary Grider (LANL)**