Department of Defense

High Performance Computing Modernization Program

Technology Insertion Process: Performance Characterization, Modeling, and Benchmarking

Mr. Cray Henry, Director

May 6, 2003

http://www.hpcmo.hpc.mil

Agenda

Benchmarking Has Real Impact

- DoD HPC Modernization Program
- Technology Insertion Process
  - Requirements
  - Selection Criteria
  - Benchmark (Brief intro — more by others later)
- Price/Performance/Outcomes

Used to Support HPCMP Acquisitions in 2001, 2002, 2003 and now 2004


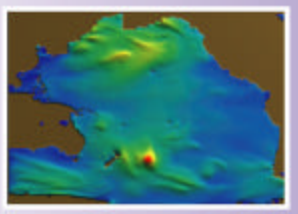

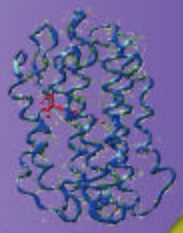

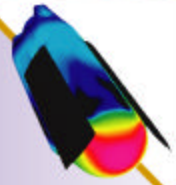

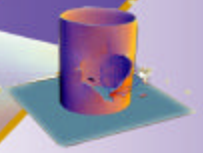


Mission

Deliver world-class commercial, high-end, high performance computational capability to the DoD's science and technology (S&T) and test and evaluation (T&E) communities, facilitating the rapid application of advanced technology into superior warfighting capabilities.

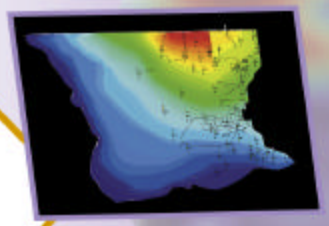
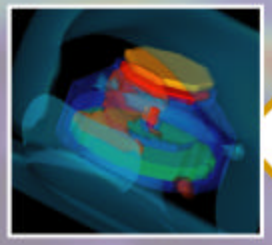

Vision

A pervasive culture existing among DoD's scientists and engineers where they routinely use advanced computational environments to solve the most demanding problems.



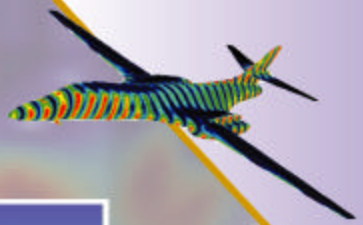


HPC Modernization Program

Supporting the Warfighter

DEPARTMENT OF DEFENSE · UNITED STATES OF AMERICA

HPC Modernization Program

Software Applications Support

CHSSI

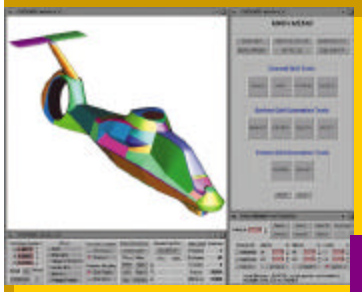

Air Force SAS
15 CHSSI Projects, 3 CTA Leaders, & 2 Portfolio Leaders

Army SAS
16 CHSSI Projects & 4 CTA Leaders & 5 Portfolio Leaders

Navy SAS
18 CHSSI Projects & 2 CTA Leaders& 1 Portfolio Leader

Defense Agencies SAS
1 CHSSI

PET

Warfighter Support

HPC Centers

Air Force HPC Centers
ASC MSRC
AEDC, AFRL/IF, SIMAF, & MHPCC DCS
1,266 Users

Army HPC Centers
ARL & ERDC MSRC
AHPCRC, SMDC, & WSMR DCs
1,355 Users

Navy HPC Centers
NAVO MSRC
FNMOC, NAWCAD & NRL DCs
1,437 Users

Defense Agencies
DARPA, DTRA, JNIC, MDA, & OTE
285 Users

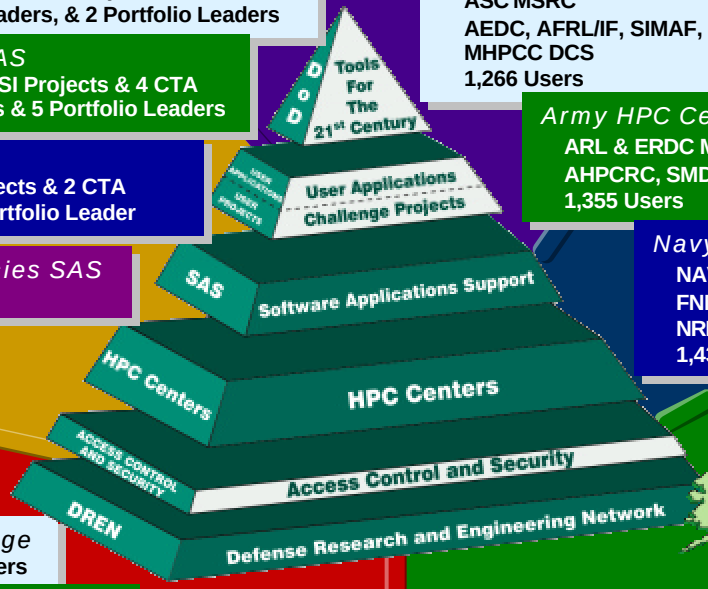

Resource Mgmt

Air Force Challenge
10 Proj. and 1,286 Users

Army Challenge
11 Proj. and 1,335 Users

Navy Challenge
16 Proj. and 1,437 Users

Army & Navy Challenge
1 Proj.

Defense Agencies Challenge
1 Proj.

Requirements & Allocations

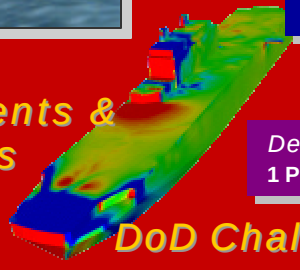

DoD Challenges

DREN

Air Force DREN
12 Sites

Army DREN
25 Sites

Navy DREN
25 Sites

Defense Agencies DREN
7 Sites

As of: August 2002

FY 2003 User Base and Requirements

- 680+ projects and 4,343 users
- Requirements categorized in 10 Computational Technology Areas (CTA)
- FY 2003 non-real-time requirements of 76 teraFLOPS-years

CSM – 544 Users

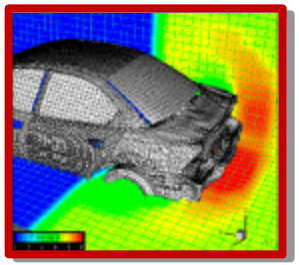

CWO – 279 Users

FMS – 290 Users

CFD – 1,111 Users

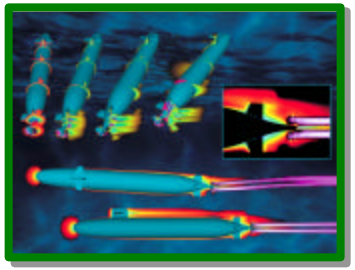

CCM – 251 Users

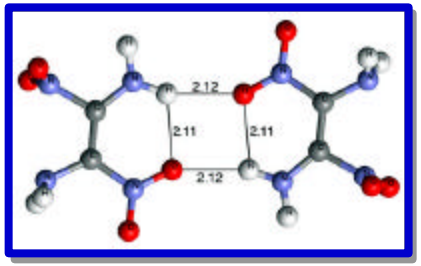

CEN – 36 Users

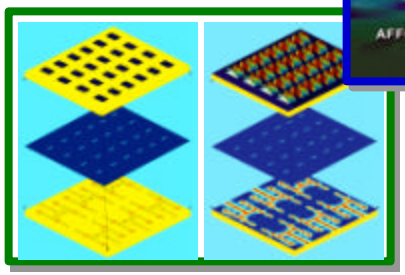

CEA – 325 Users

SIP – 311 Users

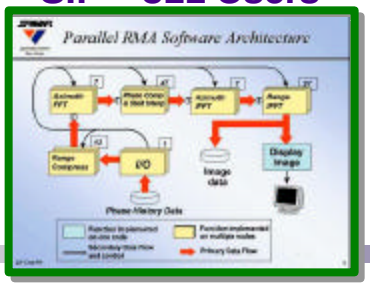

IMT – 936 Users

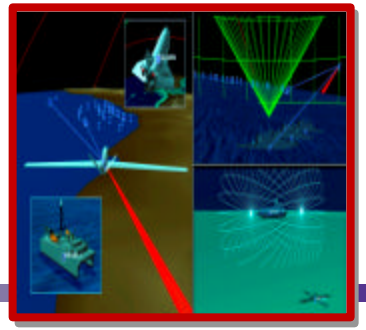

EQM – 200 Users

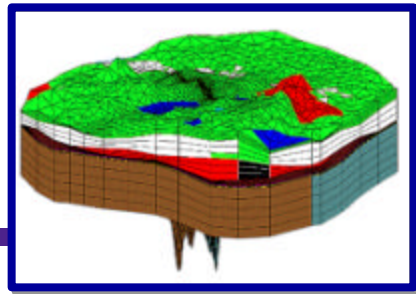

Technology Insertion (TI) Flow Chart

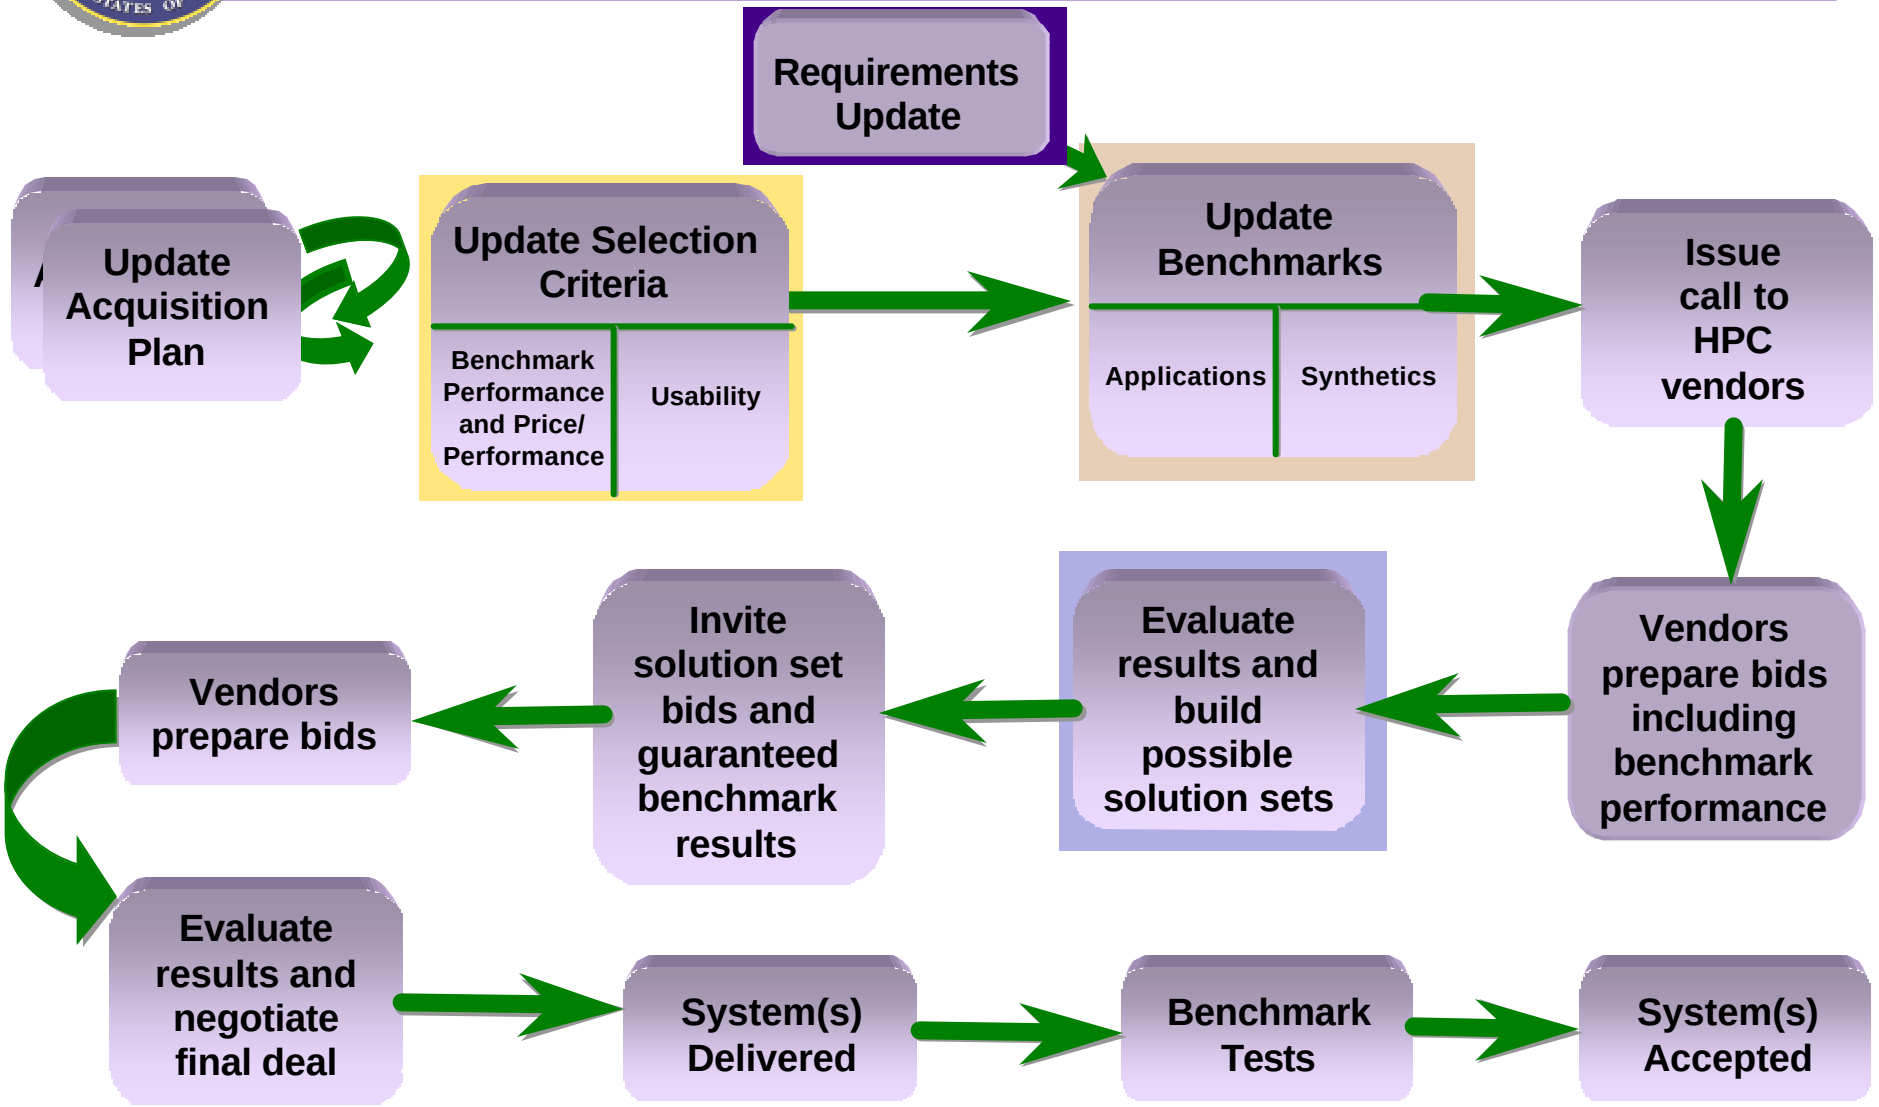

Requirements Analysis Process

Requirements Update

Objective:

- Ensure accurate HPC requirements are documented in a timely manner to impact program planning decisions
- Conduct a thorough and rigorous annual requirements analysis process to ensure accuracy

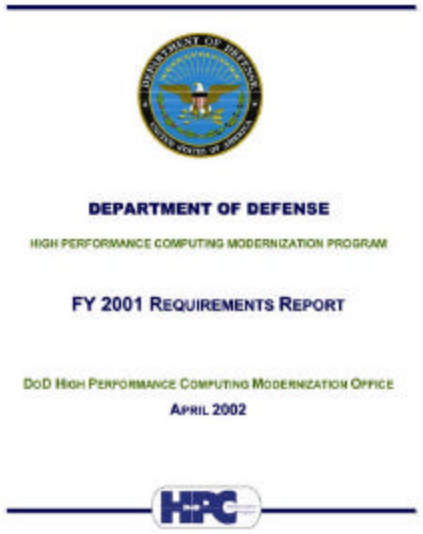

Process:

- Questionnaire - probes all aspects of HPC requirements
- Interviews - allow face-to-face clarification of detailed requirements
- Service validation - ensures that only approved/funded projects included
- Requirements analysis database - detailed profile of user base and its requirements

A user must be associated with at least one validated HPC project that has been documented in the HPCMPO's requirements database

Update Selection Criteria Qualitative

- Establish Qualitative Selection Criteria
  - Financial
  - Business Model
    - Technical roadmap
    - Strategic Vision for HPC
    - Leadership and innovation in HPC technology
  - Past Performance
    - Meeting Delivery Schedules
    - MTBF
  - Maintenance/Warranties
  - Support Model
    - Problem Resolution Process
    - Service Strategies
  - Usability
    - User Perspective:
      - Programming environment
      - Parallel file system
    - Operator Perspective:
      - HW/SW Resilience/maintainability
      - Job scheduling capability/features
      - Facilities requirements

Update Selection Criteria Quantitative

- Select Quantitative Evaluation Areas
  - Performance (quickest time to solution)
  - Price/Performance (capacity)
- Select Benchmarks
  - Application – Application/Machine Performance
    - Check Usage History by # of Users
    - Check Usage History by # CPU hours used
    - Consultation with Experts
  - Synthetic – Basic Machine Performance

Benchmark Selections
Software Application Requirements

Applications Code	Number of Users			Number of Projects		
	Unclass	Class	Total	Unclass	Class	Total
Abaqus	129	1	130	32	1	33
CTH	*78*	*31*	*109*	*14*	*3*	*17*
Cobalt	*64*	*2*	*66*	*21*	*1*	*22*
NASTRAN	65	0	65	14	0	14
ANSYS	60	1	61	16	1	17
Allegra	30	25	55	3	2	5
CFD++	41	10	51	12	2	14
PARADYN	23	25	48	2	2	4
PRONTO	20	25	45	1	2	3
GASP	29	11	40	13	4	17

Benchmark Code

Benchmark Selections
Software Requirements by CPU Hours

Application Code	CPU Hours
BRL-CAD	2,750,000
MUVES	6,750,000
BEAMS	1,500,000
CTH	13,100,000
XPATCH	2,500,000
NXAIR	210,000
WIND	1,474,380
GASP	788,750
COBALT	4,627,025
GAMESS	362,850
GAUSSIAN98	882,350
ALLEGRA	478,000
ABAQUS	1,583,206
GRIDGEN	445,875

Benchmark Selections
Consultation with Application Experts

- Candidate codes discussed with
  - CTA leaders
  - Code Developers
  - Benchmark Team

Application Benchmark Codes

- CTH (CSM)
- GAMESS (CCM)
- NLOM (CWO)
- LESlie3D (CFD)
- $Cobalt_{60}$ (CFD)
- AERO (CFD)
- NAMD (CCM)

Synthetic Benchmark Codes

- **I/O Tests**
  - Generates a 2.5 TB file, fragments the file system, then runs multiple stream tests
- **Operating System Tests**
  - Measures the performance of system calls, interprocess communication, and TCP scalability
- **Memory Tests**
  - Measures memory hierarchy performance, such as memory bandwidth
- **Network Tests**
  - A set of five MPI tests (point-to-point, broadcast, allreduce)
- **CPU Tests**
  - Multiple fundamental computation kernels, BLAS routines, and ScaLapack routines

Update Benchmarks

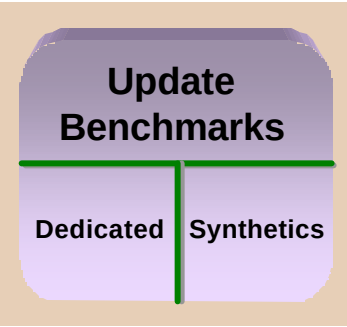

Determine workload percentages by CTA

Partition CTA percentages among benchmarks

Determine performance scores for each system on each benchmark test case at desired system size

Optimize price/performance score for each specific alternative set of systems by apportioning workload among benchmark test cases

Determine price/performance score for each viable alternative and use as major criterion for final selection

Dimensions of HPC Requirements

- Primary requirements input to acquisition process is through selection and weighting of benchmark applications
  - Suite of application codes mimics total HPCMP workload
  - Weighting of application codes timings in constructing overall performance metrics reflects size of requirements in each CTA
- Additional dimensions of requirements are checked for consistency of proposed upgrades with totality of HPCMP requirements
  - Unclassified vs. classified requirements

Percentage of Unclassified Non-Real-Time Requirements, Usage, and Allocations

Determine workload percentages by CTA

CTA	Requirements Percentage *FY [2002] (2003)*	Usage Percentage *FY 2002*	Allocation Percentage *FY 2003*	Average (25% FY 2003 Req, 25% FY 2002 Usage, 50% FY 2003 Alloc) *FY [2002] (2003)*	
CFD	[35.5%] (36.9%)	48.3%	40.7%	[43.3%] (41.6%)	
CCM	[15.5%] (18.6%)	16.4%	14.2%	[14.2%] (15.9%]	
CWO	[21.9%] (19.2%)	21.3%	21.9%	[23.3%] (21.1%)	*93.6%*
CEA	[4.1%] (4.0%)	5.1%	8.2%	[4.9%] (6.4%)	
CSM	[11.4%] (11.8%)	3.5%	9.6%	[8.3%] (8.6%)	
EQM	[3.0%] (3.2%)	0.6%	4.0%	[2.3%] (3.0%)	
SIP	[1.0%] (1.4%)	1.2%	0.2%	[0.4%] (0.7%)	
CEN	[0.5%] (0.4%)	1.3%	0.1%	[1.4%] (0.5%)	
IMT	[2.9%] (0.8%)	2.1%	0.7%	[0.9%] (1.1%)	
Other	[1.3%] (1.2%)	0.1%	0.2%	[0.4%] (0.4%)	
FMS	[2.9%] (2.6%)	0.2%	0.2%	[0.7%] (0.8%)	

FY 2003 Unclassified Non-Real-Time Job Size Requirements by CTA

Determine workload percentages by CTA

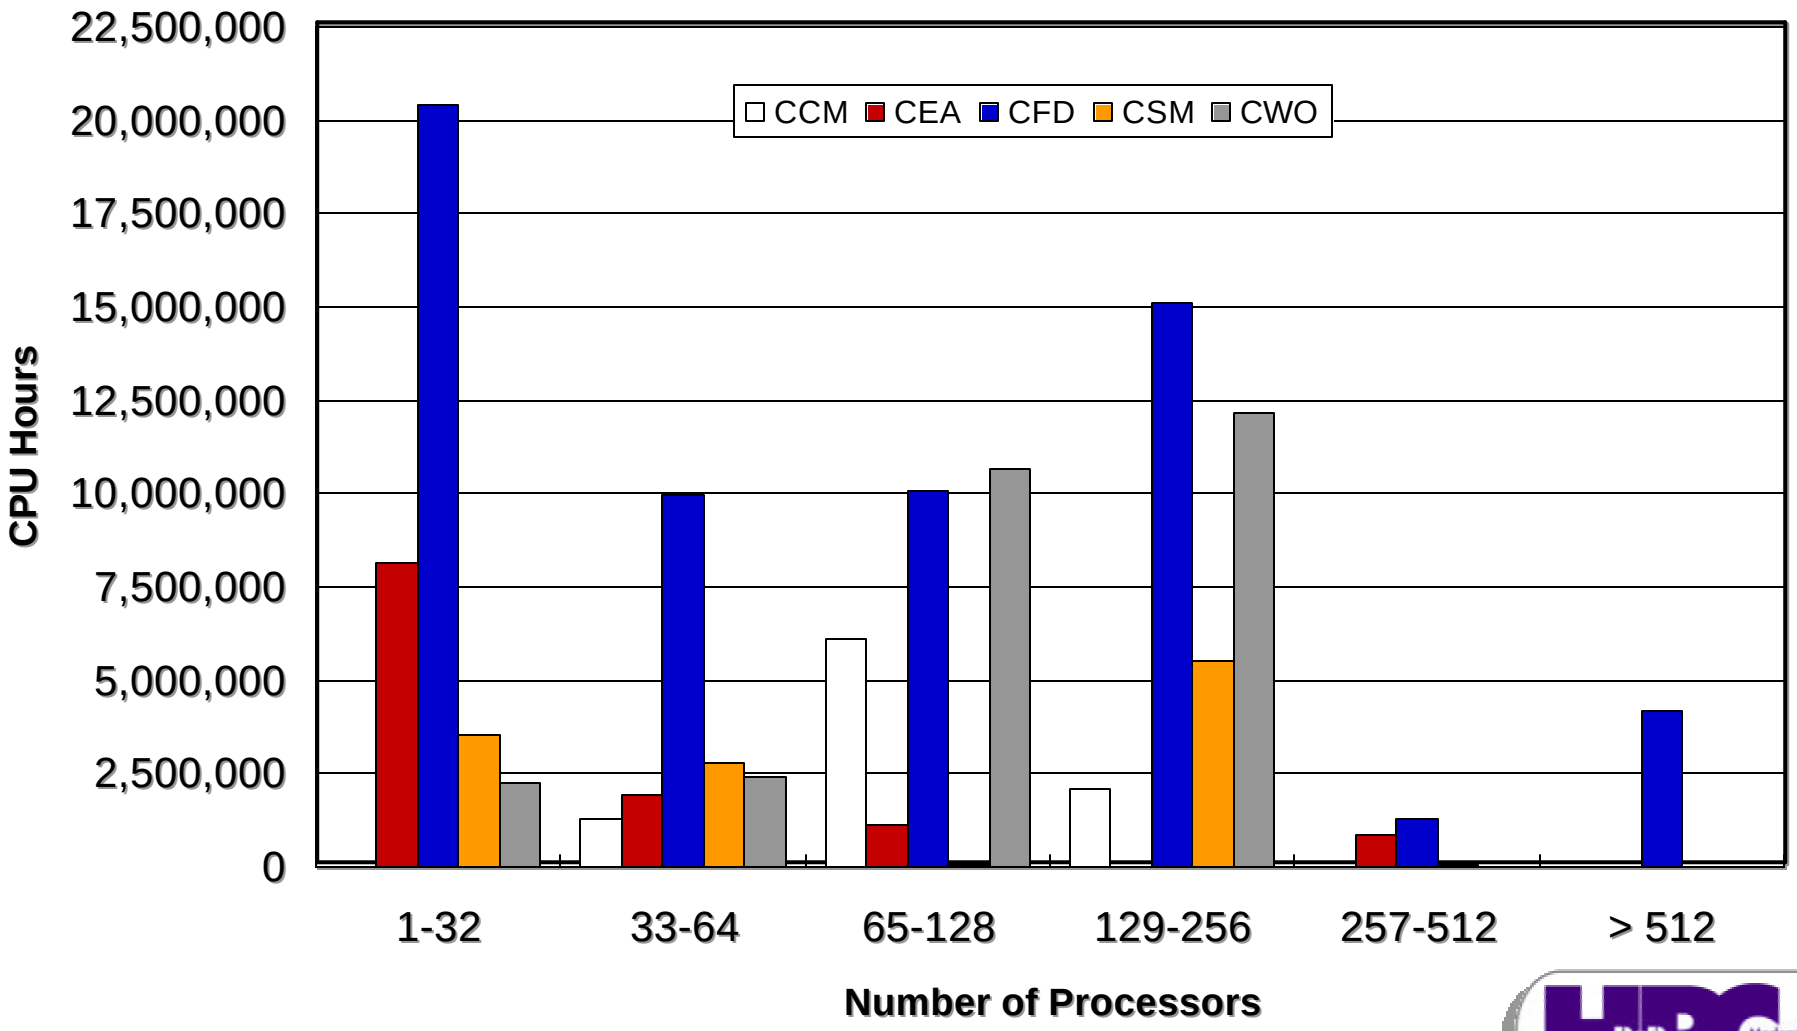

FY 2003 Percentages of Standard and Large Requirements by CTA (<=64 processors is standard)

Partition CTA percentages among benchmarks

CTA	Unclassified	
	Standard	Large
CSM	66.1%	33.9%
CFD	46.9%	53.1%
CCM	19.7%	80.3%
CEA	77.8%	22.2%
CWO	16.7%	83.3%
SIP	64.4%	35.6%
FMS	100.0%	0.0%
EQM	72.3%	27.7%
CEN	23.4%	76.6%
IMT	82.9%	17.1%
Other	90.8%	9.2%
Overall	41.1%	58.9%

Benchmark Suite

Partition CTA percentages among benchmarks

Dedicated: (80%)

- CTH (CSM)
  - Large xx%
  - Small yy%
- GAMESS (CCM)
  - Large …
  - Small …
- NLOM (CWO)
  - Large …
  - Small …
- LESlie3D (CFD)
  - Large …
  - Small …
- $Cobalt_{60}$ (CFD)
  - Large …
  - Small …
- AERO (CFD)
  - Large …
  - Small …
- NAMD (CCM)
  - Large …
  - Small …

Synthetic (20%)

- I/O Tests xx%
- Operating System Tests yy%
- Memory Tests …
- Network Tests …
- CPU Tests …

Emphasis on Performance
Time to Solution

Determine performance scores for each system on each benchmark test case at desired system size

- **Establish a DoD standard benchmark time for each application benchmark case**
  - **NAVO IBM SP P3 chosen as standard DoD system**
- **Benchmark timings (at least three on each test case) are requested for systems that meet or beat the DoD standard benchmark times by at least a factor of two (preferably four)**
- **Benchmark timings may be extrapolated provided they are guaranteed, but at least one actual timing must be provided**

CTH Large

NAVO IBM SP P4 — 1124 Processors

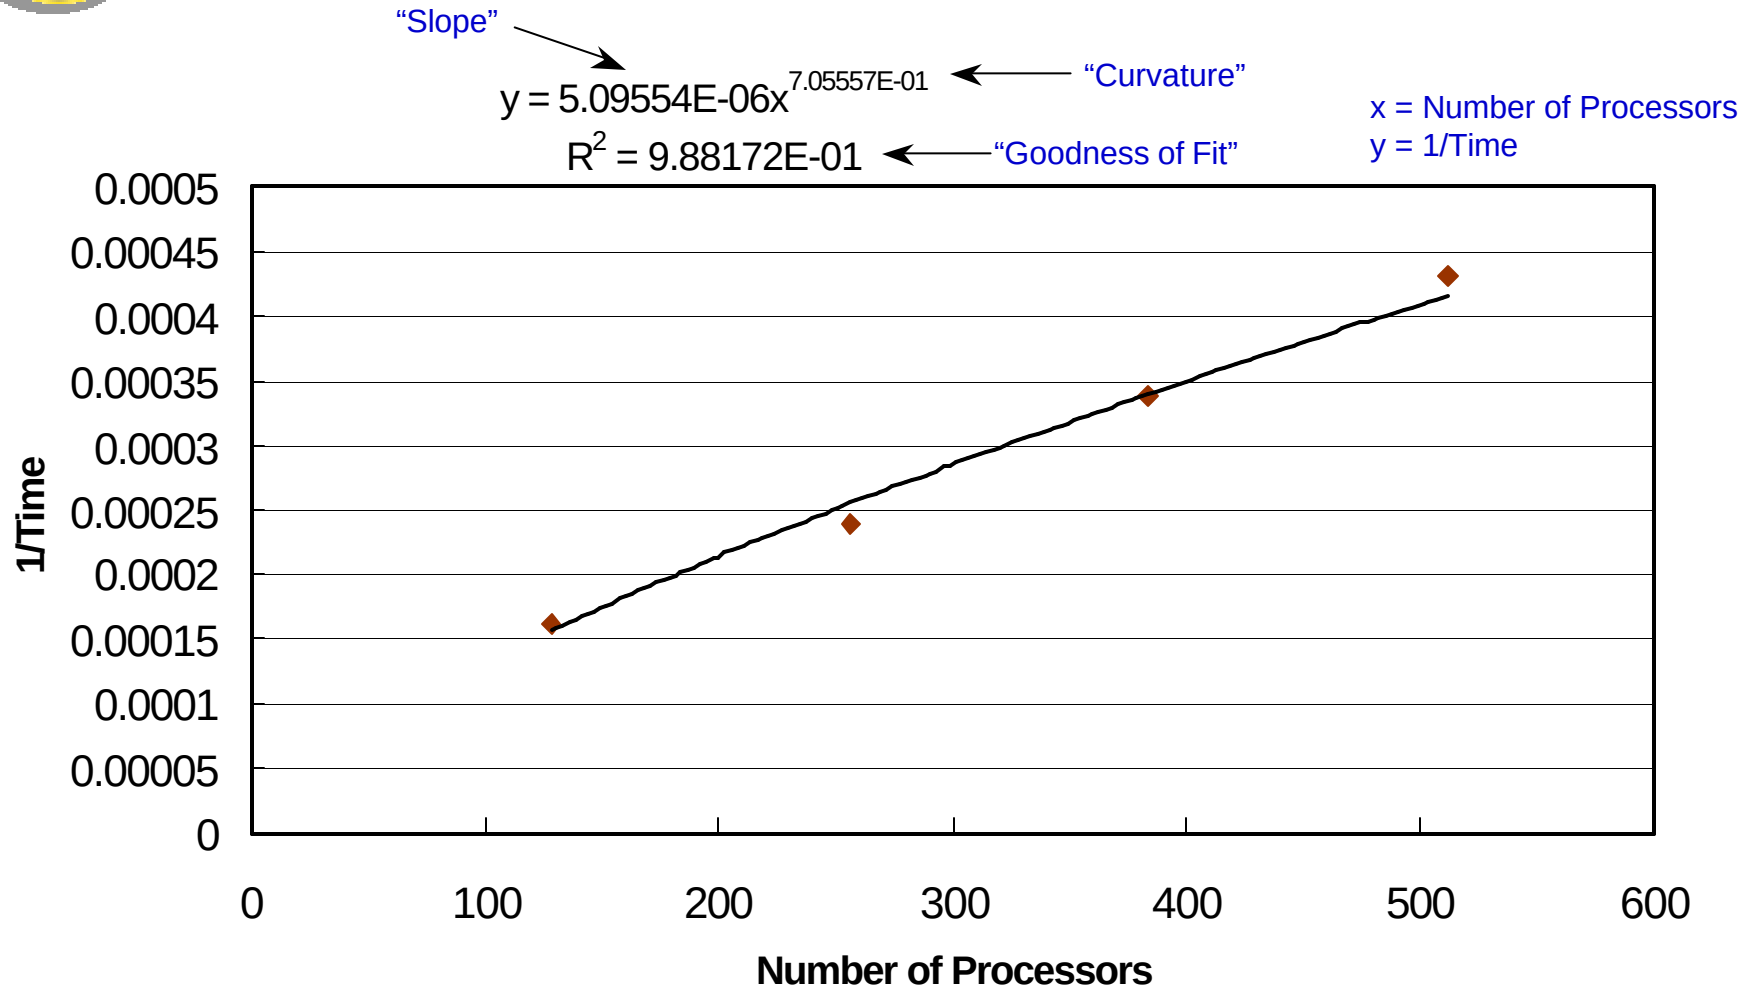

CTH Large

NAVO IBM POWER4 — 1124 Processors

DoD Standard System Peformance			
256	7245	← Seconds	
System Performance at Required Size			
281	3673.74	0.0002722	
Raw Performance =		1.972	
Sys Size to Match Perf		107.3	
Capacity Performance =		2.618	

Number of processors

1/4 of total system size for large test cases

Projected time at required system size

Power law performance at required system size

Ratio of 256-processors of standard system time to projected time at required system size

Ratio of required system size to number of processors required to match standard system performance

Number of processors needed to match standard system performance based on power law fit

Result: Performance Score for Each Offered System and Benchmark

Evaluate results and build possible solution sets

System	# Proc	CTH Std 6.53%	CTH Lg 3/35%	Aero 11.94%	Cobalt S 11.20%	Cobalt L 12.68%
Cray X1	128	x.xx	x.xx	x.xx	x.xx	x.xx
Cray X1	64	x.xx	x.xx	x.xx	x.xx	x.xx
HP SC45 1.0GHz	188	x.xx	x.xx	x.xx	x.xx	x.xx
IBM Pow 4 1.7GHz 690	160	x.xx	x.xx	x.xx	x.xx	x.xx
IBM Pow 4 1.7GHz 690	128	x.xx	x.xx	x.xx	x.xx	x.xx
LN Pent 4 2.4GHz Q	512	x.xx	x.xx	x.xx	x.xx	x.xx
LN Pent 4 2.4GHz Q	256	x.xx	x.xx	x.xx	x.xx	x.xx
LN Pent 4 2.4GHz M	512	x.xx	x.xx	x.xx	x.xx	x.xx
LN Pent 4 2.4GHz M	256	x.xx	x.xx	x.xx	x.xx	x.xx
SGI O3000 600MHz	256	x.xx	x.xx	x.xx	x.xx	x.xx
SGI O3000 700MHz	1024	x.xx	x.xx	x.xx	x.xx	x.xx
SGI O3000 700MHz	512	x.xx	x.xx	x.xx	x.xx	x.xx

HPC System Performance Results

Normalized Capability Performance Scores

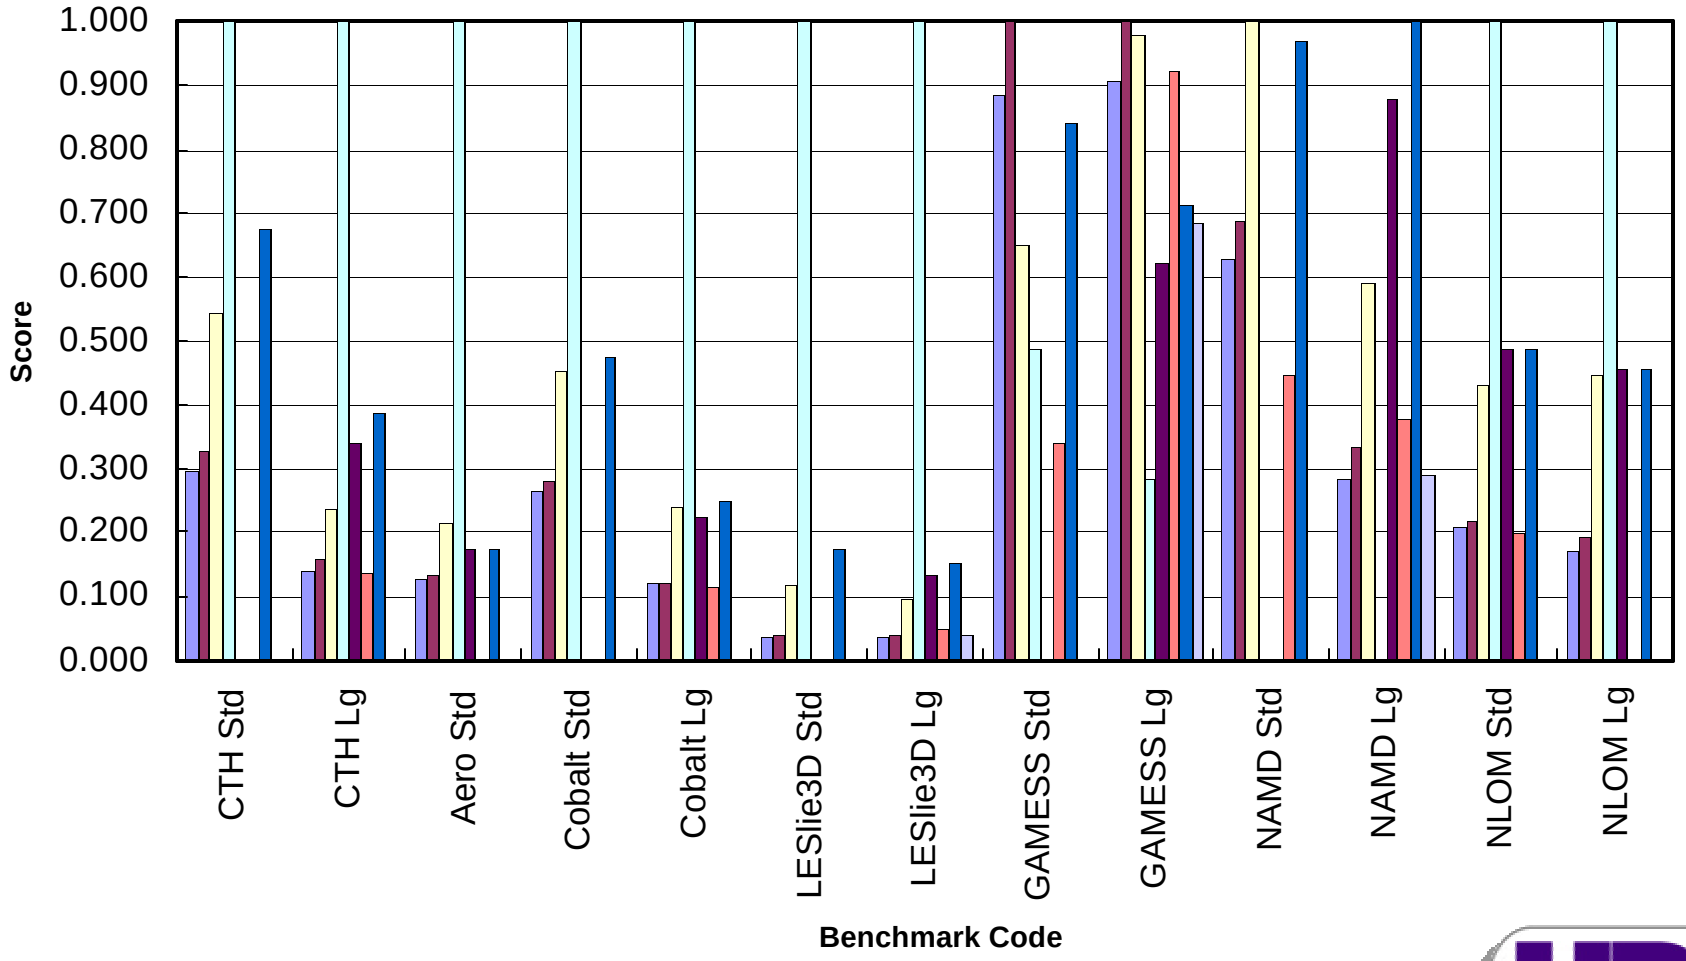

Solution Set Building

Evaluate results and build possible solution sets

System					CTH Std	CTH Lg	Aero	Cobalt S	Cobalt L
Unclassified Benchmark Weights =					6.53%	3/35%	11.94%	11.20%	12.68%
Classified Benchmark Weights =					XX	XX	XX	XX	XX
System	# Proc	Number	Cost($M)	Total					
Cray X1	128	0	$1	$0.000	0.000	0.000	0.000	0.000	0.000
Cray X1	64	0	$1	$0.000	0.000	0.000	0.000	0.000	0.000
HP SC45 1.0GHz	188	0	$1	$0.000	0.000	0.000	0.000	0.000	0.000
IBM Pow 4 1.7GHz 655	512	0	$1	$0.000	0.000	0.000	0.000	0.000	0.000
IBM Pow 4 1.7GHz 690	160	0	$1	$0.000	0.000	0.000	0.000	0.000	0.000
IBM Pow 4 1.7GHz 690	128	0	$1	$0.000	0.000	0.000	0.000	0.000	0.000
LN Pent 4 2.4GHz Q	512	0	$1	$0.000	0.000	0.000	0.000	0.000	0.000
LN Pent 4 2.4GHz Q	256	0	$1	$0.000	0.000	0.000	0.000	0.000	0.000
LN Pent 4 2.4GHz M	512	0	$1	$0.000	0.000	0.000	0.000	0.000	0.000
LN Pent 4 2.4GHz M	256	0	$1	$0.000	0.000	0.000	0.000	0.000	0.000
SGI O3000 600MHz	256	0	$1	$0.000	0.000	0.000	0.000	0.000	0.000
SGI O3000 700MHz	1024	0	$1	$0.000	0.000	0.000	0.000	0.000	0.000
SGI O3000 700MHz	512	0	$1	$0.000	0.000	0.000	0.000	0.000	0.000
SGI O3000 700MHz	256	0	$1	$0.000	0.000	0.000	0.000	0.000	0.000
SGI O3000 700MHz T	256	0	$1	$0.000					
					0.000	0.000	0.000	0.000	0.000
Total for Alternative				$0.000	0.000	0.000	0.000	0.000	0.000
Application Percentage					0.000%	0.000%	0.000%	0.000%	0.000%

Total Performance Score

Benchmarks

Today

Dedicated Applications

- 80% weight
- Real codes
- Representative data sets

Synthetic Benchmarks

- 20% weight
- Future look
- Focus on key machine features

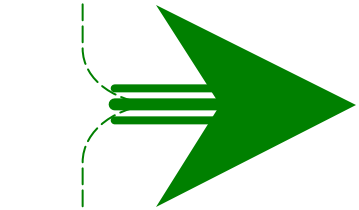

Tomorrow

Synthetic Benchmarks

- 100% weight
- Coordinated to application "signature"
- Performance on real codes accurately predicted from synthetic benchmark results
- Supported by genuine "signature" databases

Next 2–3 year key — must prove that synthetics benchmarks and application "signatures" can be coordinated

Summary

Benchmarking Has Real Impact

- $120M in decisions over last 3 years
- $100s of millions in decisions over the next decade

Synthetics performance coordinated to application signatures is the next huge step.

Make it Happen!