HPC User Forum 48

September 2012

Merle Giles
Business & Economic Development
NCSA Consulting

National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign

What Industry Requires to Move Science and R&D Forward

National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign

Innovation

2007, 2009, 2012

National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign

Competitive Advantage - 2007

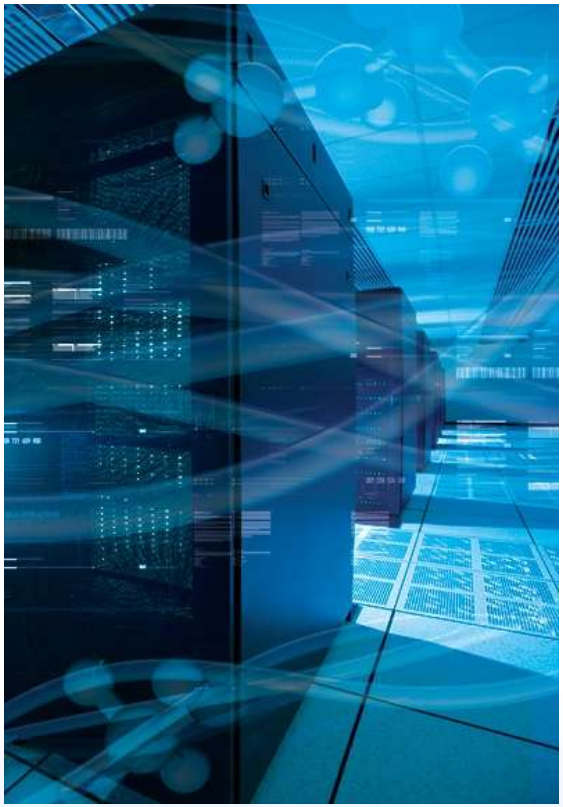

Winning globally requires HPC-driven solutions.

Suzy Tichenor, Council on Competitiveness, March 2008

U.S. FORTUNE100® 2010 Revenues ($ in billions)

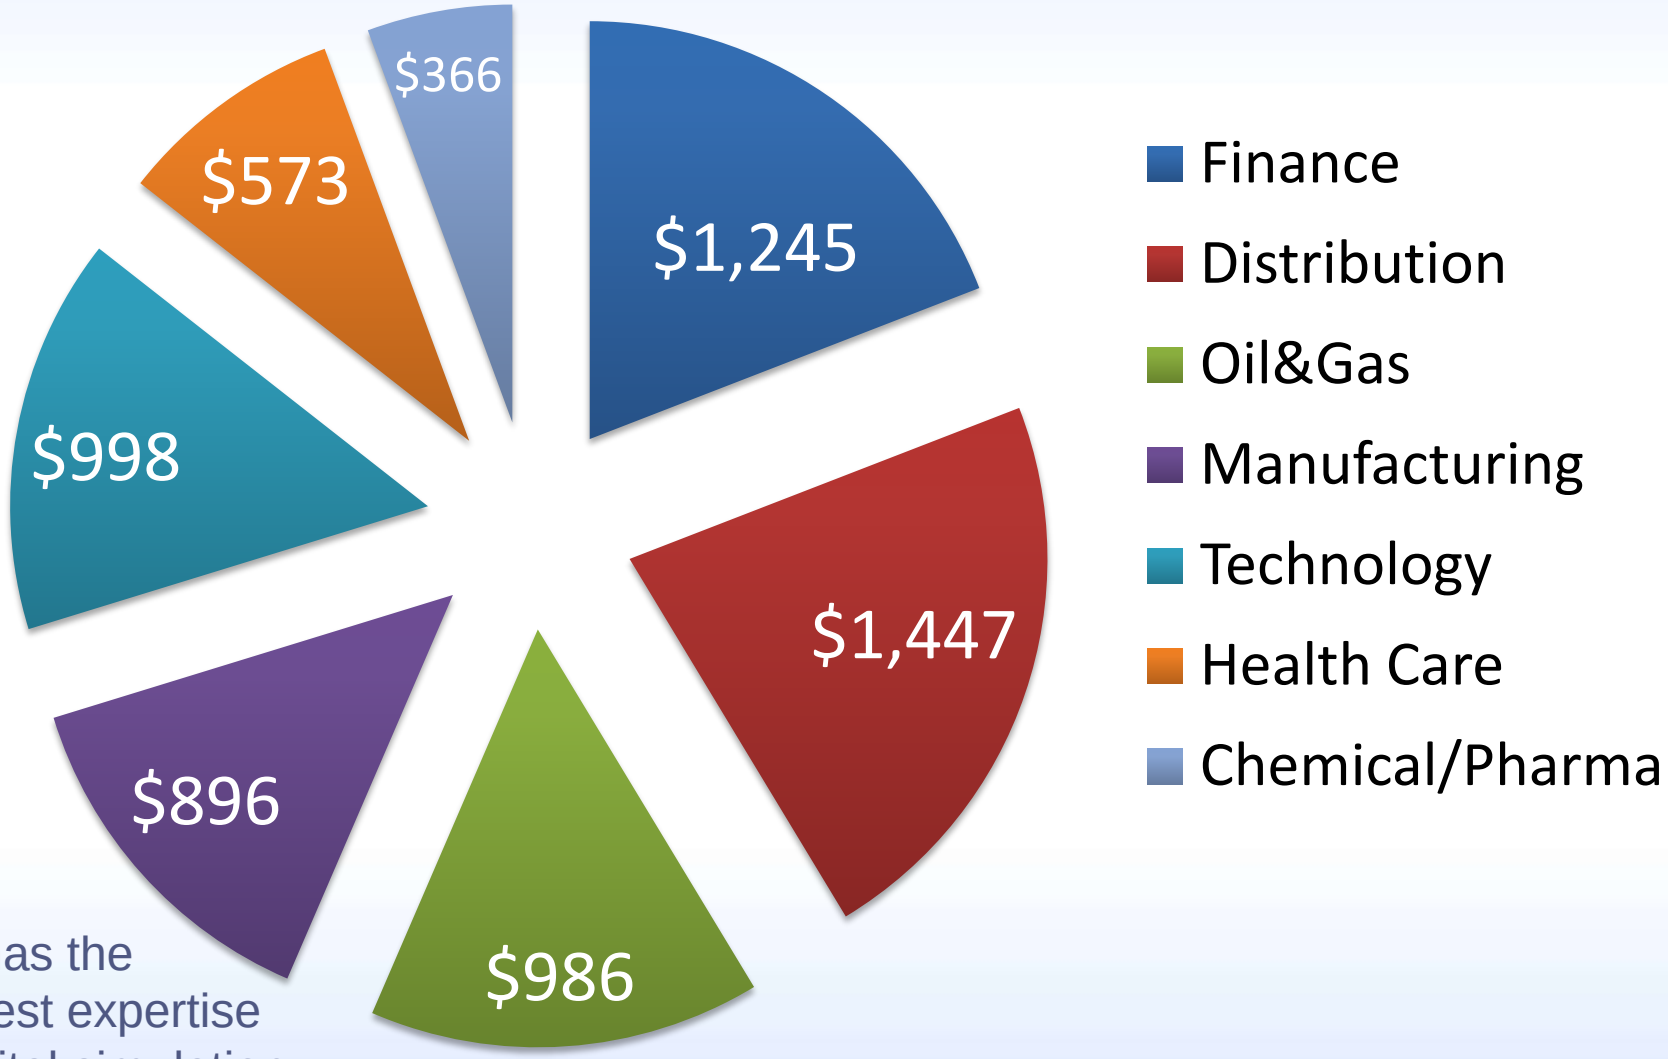

Mfg has the deepest expertise in digital simulation

Discovery is No Longer Sufficient

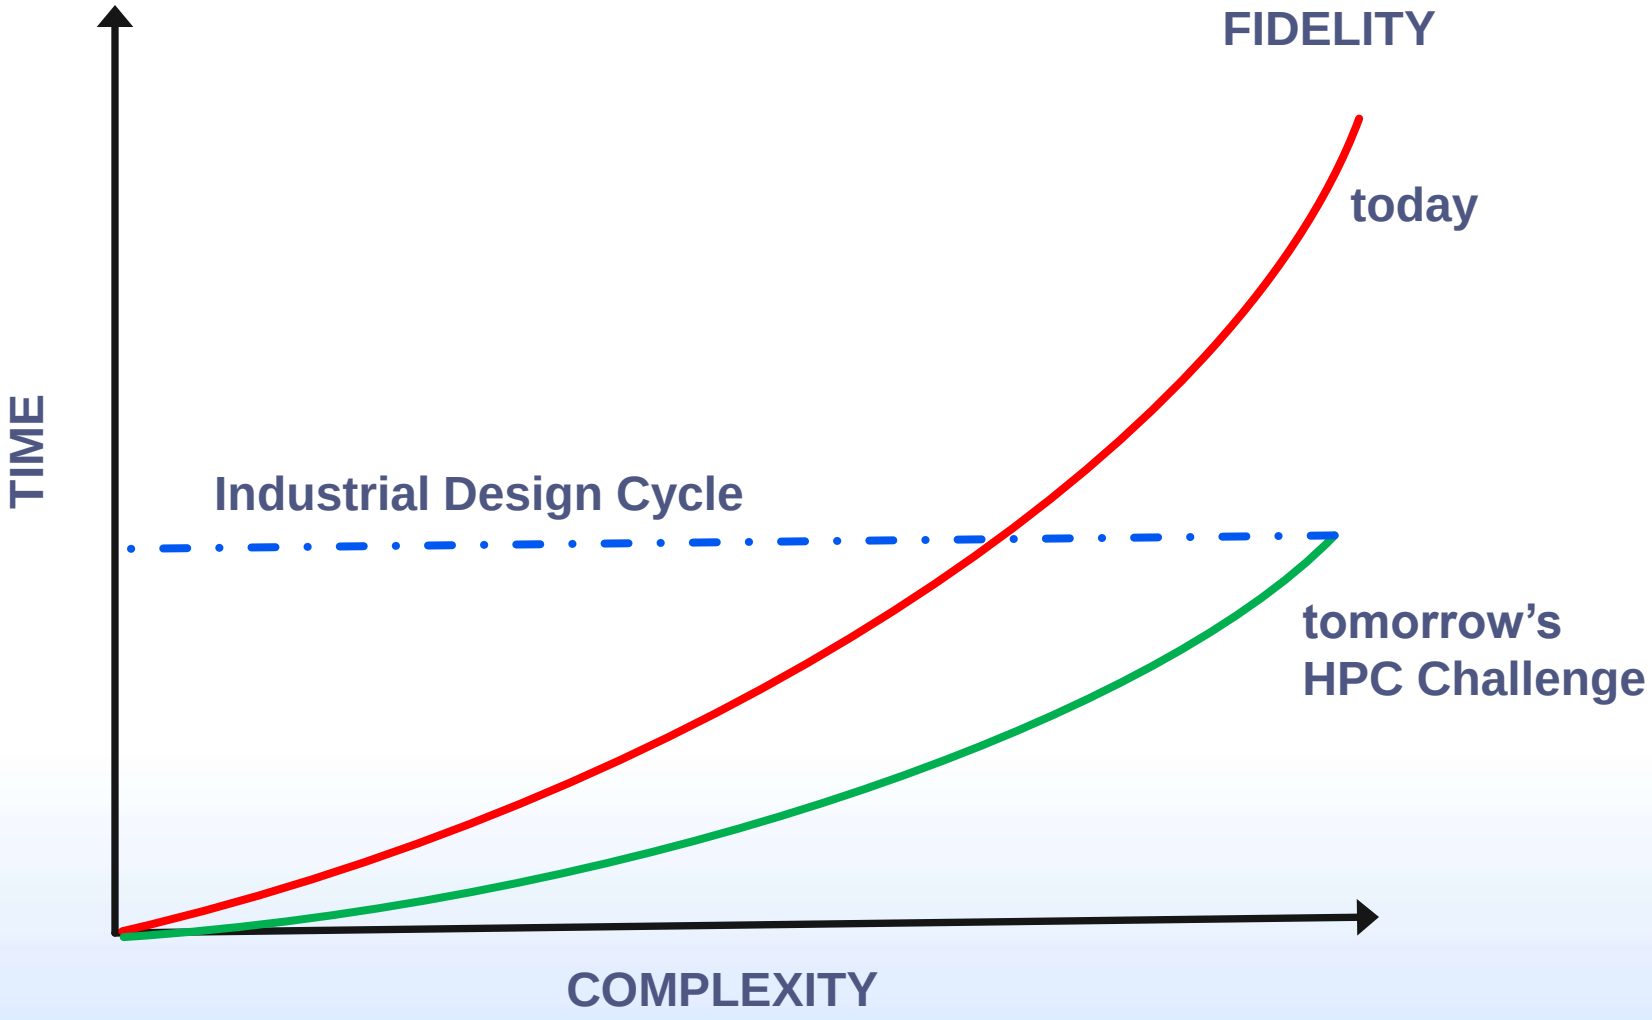

Source: Rolls-Royce Shyam Neerarambam

Innovation - 2009

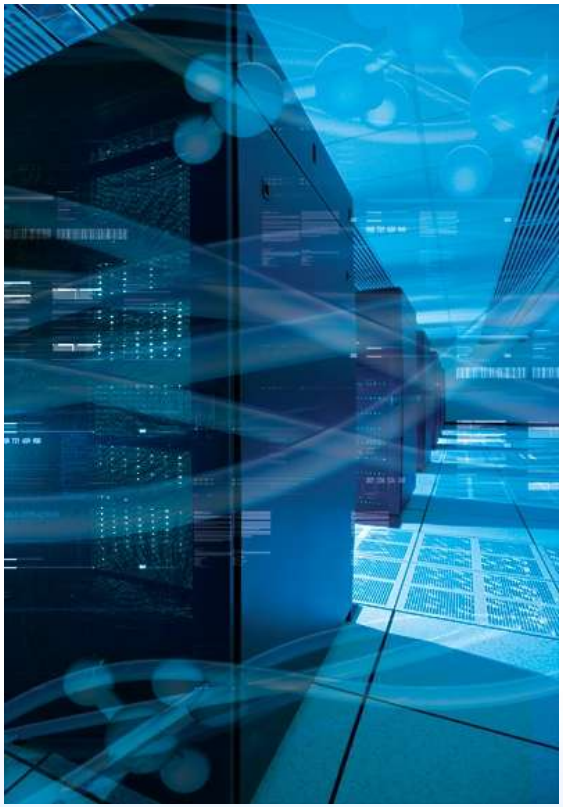

Innovation occurs at the intersection of basic and applied research.

Regina Dugan, Darpa Director, March 2009

Industry Surveys for UIUC (2012)

- Industry-related technical skills + soft skills
- Industry relevance in curriculum
- Accelerated rate of innovation
- Increased speed of research (5 years too long)
- Real-time problem solving
- Lack of more applied research will affect funding
- Thematic clusters or researchers
- Collaborate to compete w/r/t intellectual property
- International exposure for graduates
- Convert Big Data into Knowledge
- Value extraction from structured/unstructured data
- Predictive analytics


BOEING

CAT

JOHN DEERE

GE

P&G

ROLLS
ROYCE
Rolls-Royce

Illinois Rocstar

JOHN ZINK

NDEMC
Leading 21st Century Manufacturing

CRAY

HDF

Nimbis Services

DELL

Microsoft

Adaptive
COMPUTING

ExxonMobil
Upstream Research

bp

bhpbilliton

ADM

DARK ENERGY
Survey
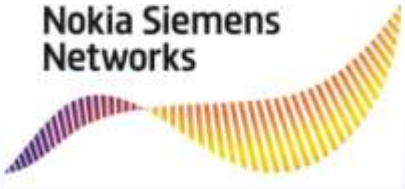
Nokia Siemens
Networks

WATERBORNE
ENVIRONMENTAL, INC.

Innovation - 2012

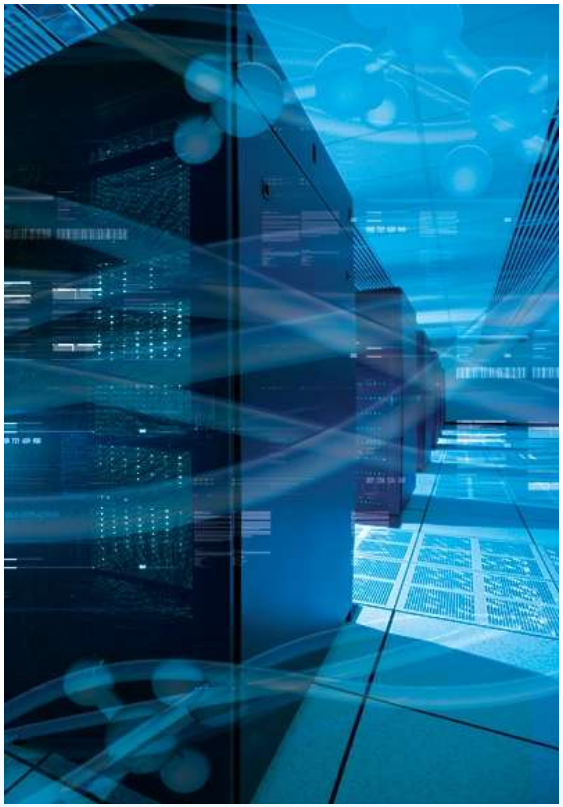

Innovation occurs at the intersection of HPC and additive manufacturing.

NNMI, Merle Giles, NCSA, September 2012

Product Lifecycle Management

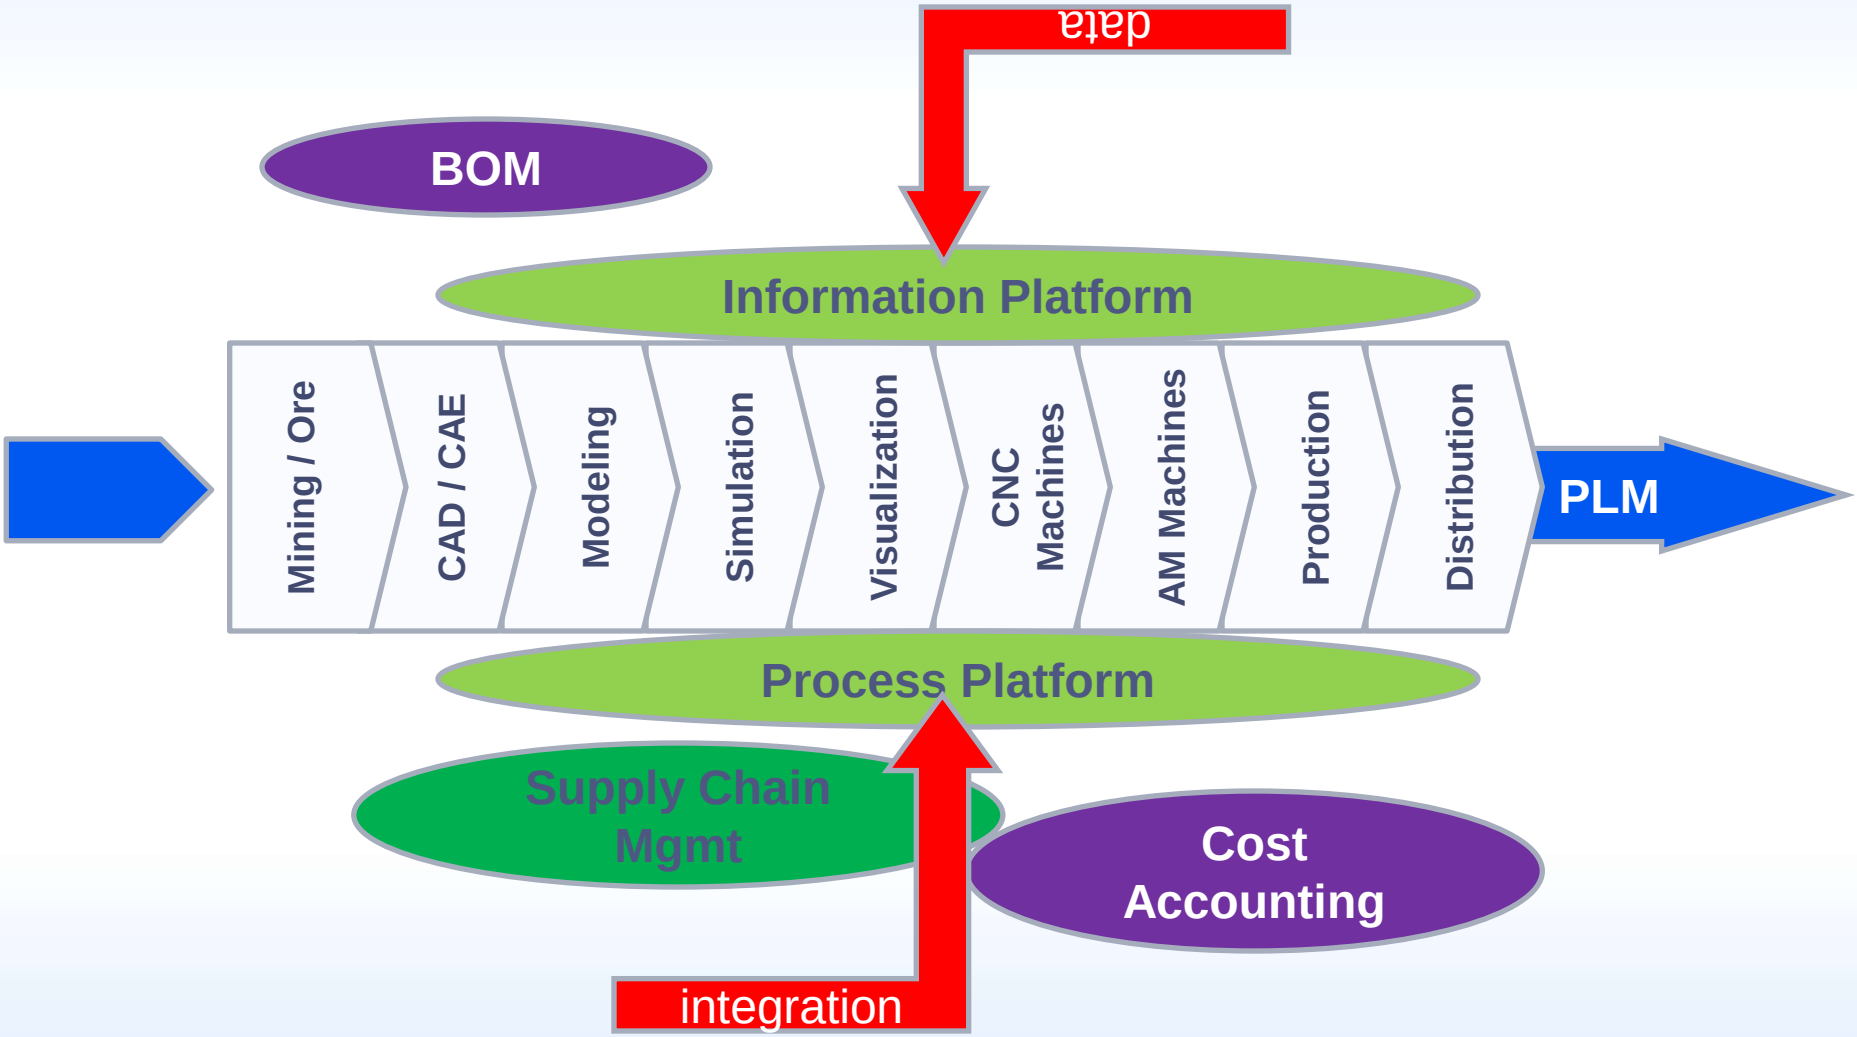

NNMI 2012

- National Network for Manufacturing Innovation
- Pilot $30M federal + $30 required match
- Focus on Additive Manufacturing
- Areas of research exist:
  - Powder base of exotic materials
  - Surface smoothness - 1000X improvement in quality
  - Integrated data (ore, design, mfg, viz, sim, lifing, . . .)
  - Remanufacturing / Resurfacing
  - Conversion of part specs to 3D digital models
  - Near-real-time simulation for on-demand AM
  - DoD want this NOW!

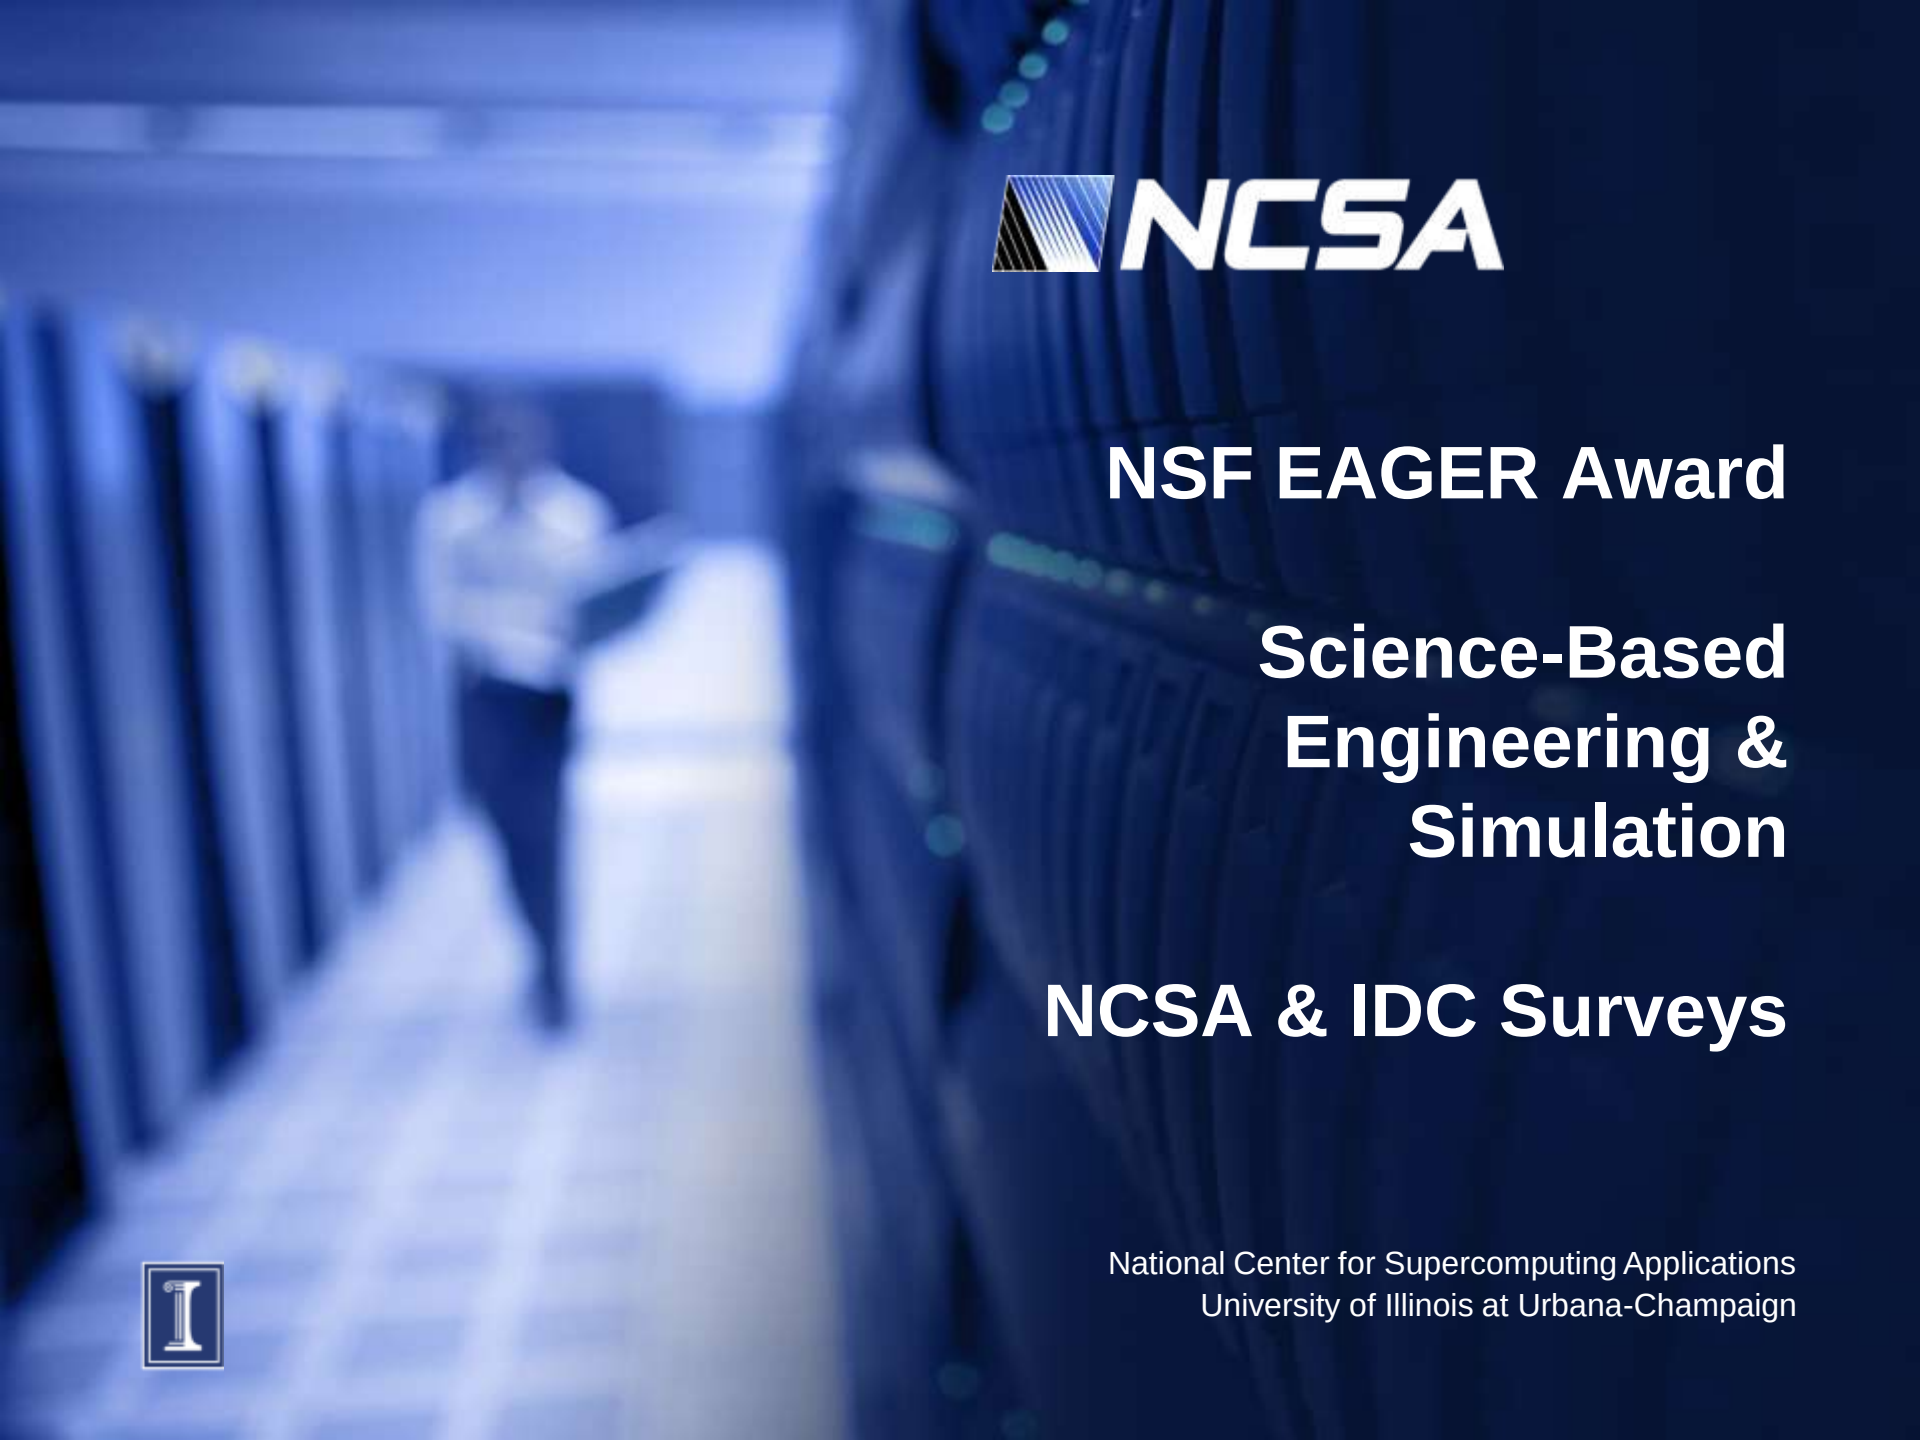

NSF EAGER Award

Science-Based Engineering & Simulation

NCSA & IDC Surveys

National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign

NDEMC

- A manufacturing/government partnership targeted at increasing the use of advanced modeling and simulation among SMEs
- Aneesh Chopra –best bottom-up business model today
- Merle Giles, NCSA – now more than 3 federal agencies can spell HPC
- Chopra – we need all hands on deck!

NDEMC – OEM Comments

- **Do NOT dumb down the product, but design only what is needed**
- Looking for a 10X payback and 10% investment in VPD
- **Too much island development today – mush integrate**
- Objectives include 5X speedup and 5X decrease in cost
- Forging machines, heat-treat, lathes all have digital interfaces – but do they talk to the digital designers?
- **Sam Allen, John Deere CEO – HPC helps accelerate companies away from poor design decisions**

NDEMC – SME Comments

- 3D model with 1 million elements takes 1.5 days
- Looking to reduce material, but only doing 2D CAD
- **We build to spec – how can we innovate?**
- MS&A is used for durability; 99% are validated with pressure; need to understand vibration, but lack detail
- **Must understand product lifing**
- This code won't do multiphysics
- We test to spec, but cannot afford to test to failure
- Quad-core machine takes 4 months to run the model

Industry Surveys for NSF (2012)

- Document which HPC applications manufacturers are using today and their future needs
- Determine the scalability of these applications
    - And the underlying science that may be limiting the scalability
    - Identify the barriers preventing greater scalability
- Determine what methods the manufacturers are using, if any, to get past scalability limits
- Recommend areas where additional investments, training and/or research could improve the scalability and usability of applications
- Surveyed 30 organizations (paper, phone, f2f)

Respondent Profiles

- 30 organizations participated:
  - 76.6% Industry
  - 16.7% Government
  - 6.7% Academia

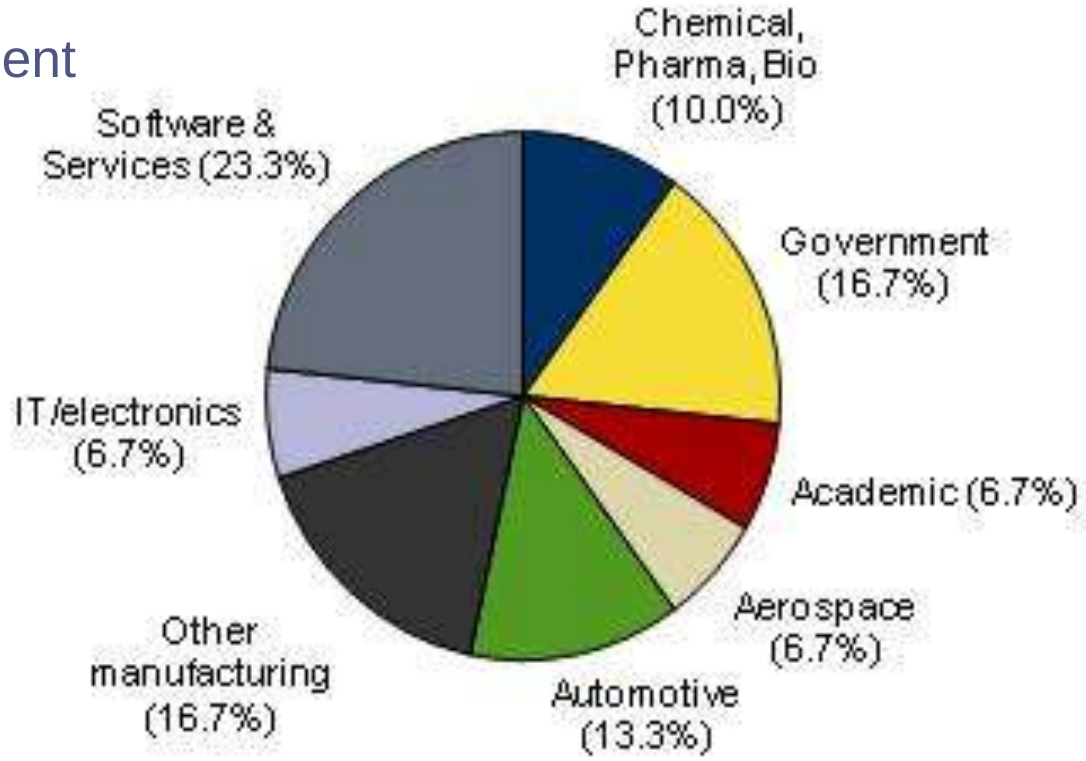

Respondent Profiles

- Organizations ranged greatly in size, from 22 to 164,000 employees
- R&D investments among the responding organizations ranged from $515 million to a whopping $5 billion
    - As a percentage of revenues, R&D expenditures varied from 3% to 45% and averaged 18.2%
    - The responding companies invested about 6.5 times more, as a percent of their revenues, than did all U.S. companies as a group

Survey Findings: Software Sources

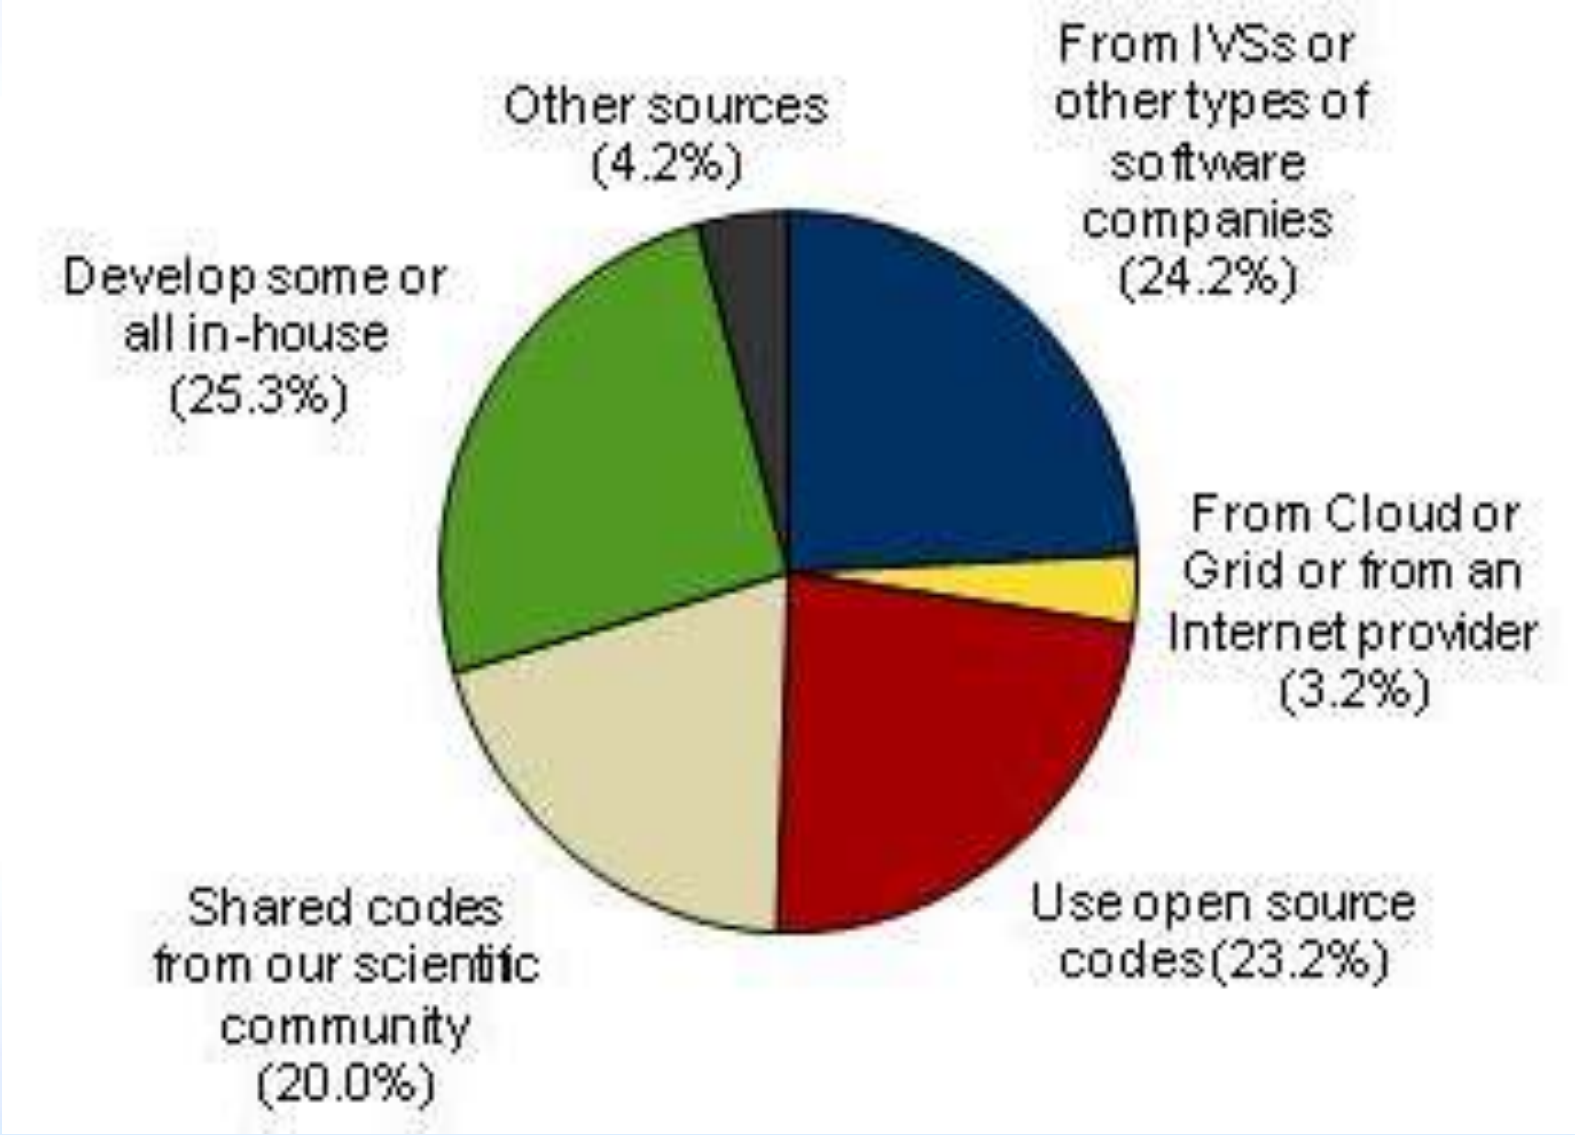

Scientific Disciplines Used In Applications

When asked to identify the scientific disciplines and sub-disciplines that underpinned their top three applications

- The participants provided a relatively long, varied list of disciplines
- Most cited disciplines were:
  - Fluid mechanics/aerodynamics, Mathematics, Physics, Bio-life sciences, Meteorology/climate science, Chemistry, Electronics/electrical theory, Structural mechanics, Combustion, Information management/statistics, Visualization
- Most cited sub-disciplines were:
  - Radiation biology, Nuclear physics, Oceanography, Proteomics, Astrophysics, Particle physics, Quantum chemistry, Computer science, Statistics

Programming Models Used

The programming models used in the top 3 applications:

- MPI was the main model for parallelism
- OpenMP was used for shared memory programming
- A Hybrid combination of MPI and OpenMP occurred often enough to point to a trend
- Only a few instances of using CUDA and OpenCL for programming GPGPUs.
- Just one mention of MapReduce/Hadoop

Longevity of Existing Applications

Q: When asked if their existing applications will meet their requirements for the next five years:

- 38.9% said the underlying mathematical model/ algorithm needs to be improved
- 31.2% said the underlying science needs to be improved
- 15.5% applications as now written would meet their requirements
- 14.5% said that a factor other than those cited above needs to be addressed

Scalability Limitations to Applications

- 31% of the surveyed organizations said that the science underlying their most important applications need to be improved as a basis for advancing the performance of their codes

- Respondents next hoped to advance one or more orders of magnitude beyond the existing limits of their scalability and resolution/realism
    - The three common factors cited include:
        - The size of the HPC resources
        - The models
        - The algorithms

Limitations To Providing Realism

Q: When asked about the main technical limitations that might be preventing the sites from dramatically increasing the problem size or other dimensionality of their simulations:

- Models/algorithms are not scalable enough
- Inadequate latency/bandwidth
- Inadequate processing power
- Underlying science needs advancement
- Software developer time would be too long
- Improved methodologies are needed

Workforce Capability – Training is Greatly Needed

- Three-quarters of the participants said their people can handle today's technologies but need retraining to handle next-generation systems.

What are the capabilities of your current work force to implement changes that improve scalability and performance of applications using latest HPC hardware and software technologies?

	Percentage Of Respondents
Workforce is capable of adequately handling the current generation technologies but need to be retrained for next generation systems	75.0%
Work force fully equipped to handle current technologies and are able to provide technology leadership for next generation systems	12.5%
Workforce is inadequate to handle the complexity of HPC and is a serious impediment to adopting advanced HPC simulation based engineering	12.5%
Workforce is capable of maintenance of current infrastructure but is inadequate to develop new solutions for current generation or future generation systems	0.0%

Source: IDC, 2012

Desired Improvements in Applications

- Most of these near-term improvements fell into the following categories:
  - Higher-resolution meshes
  - Improved mathematical models and algorithms
  - Improvements to the underlying physics
  - Better methods for data integration and analysis
- 44.4% said that the current science has the potential to provide "a large amount of additional realism”
- In a majority of cases, the known science has advanced beyond the ability of industrial applications to exploit it fully
- Additional investments to advance the known science are needed to drive industrial innovation and competitiveness forward

Where People Get Help

- The most popular source of help is relying completely on internal staff (24% of all responses)
- Closely followed by getting help from the sites' hardware vendors (21% of responses) and ISVs (19%)

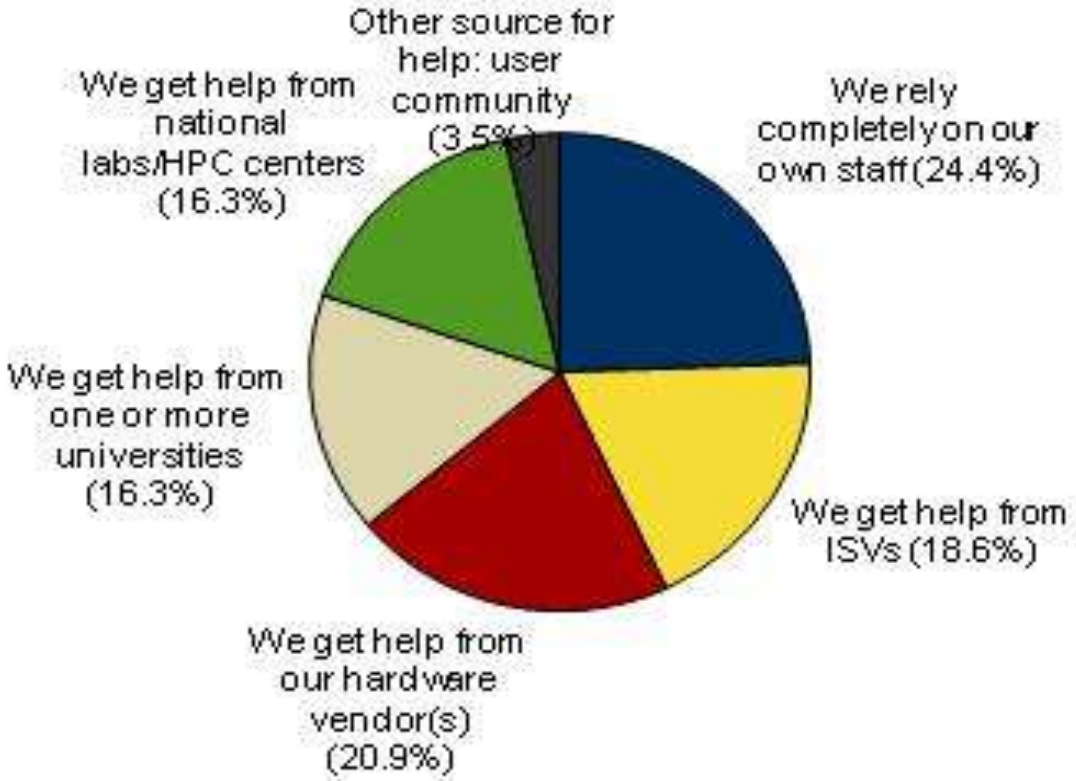

Willingness To Work With NCSA/NSF

- About 63% of the responses favored a high degree of involvement in pursuing these improvements

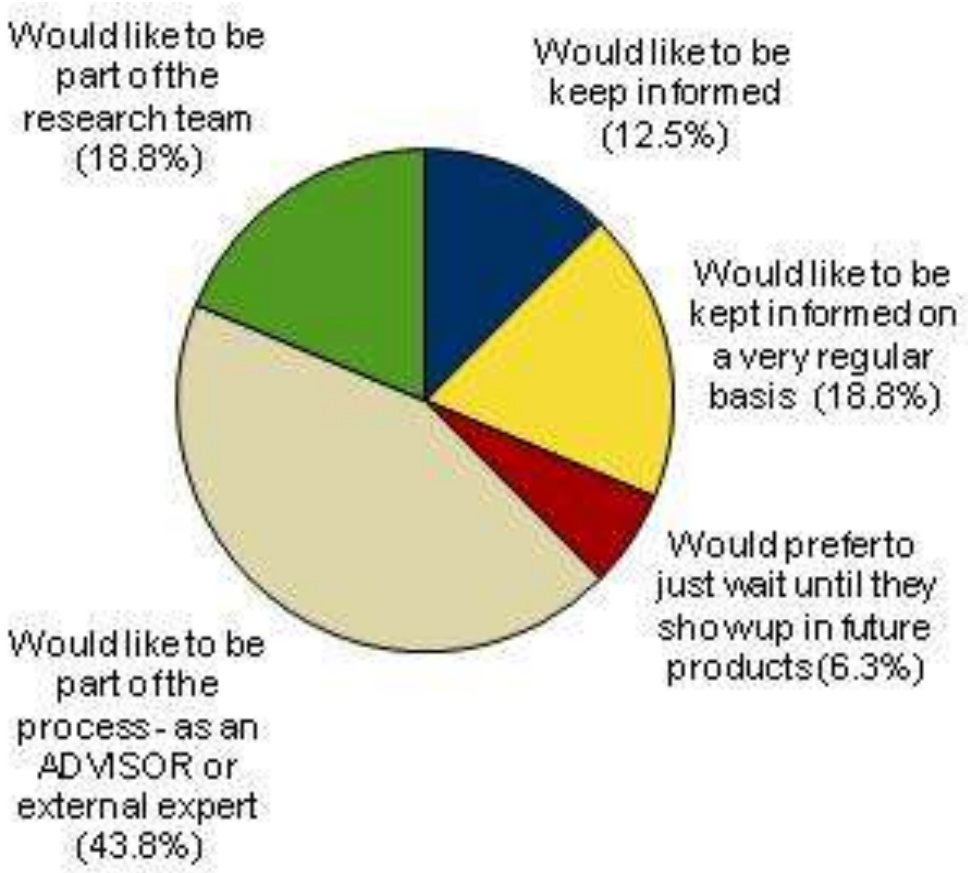

In Summary:
Areas Needing Most Improvement

- Higher-scaling models and algorithms
- More training for personnel
- Advances in the underlying science (e.g., physics, radiation biology)
- Better-performing hardware systems
- Better interconnects/networking (bandwidth, latency)
- More attractive software licensing models

In Summary

- Only 15% say that their current applications would meet their requirements for the next 5 years
  - 39% say that the underlying mathematical model/algorithm needs to be improved
  - 31% say that the underlying science needs to be improved
  - 30% of the organizations named specific ways in which they now have to "dumb down" their problems in order to complete the runs in reasonable amounts of time
- 63% of the responses favored working with NCSA/NSF in pursuing needed improvements

Scientific Disciplines Used In Applications

- When asked to identify the scientific disciplines and sub-disciplines that underpinned their top three applications
  - The participants provided a relatively long, varied list of disciplines
  - Most cited disciplines were:
    - Fluid mechanics/aerodynamics, Mathematics, Physics, Bio-life sciences, Meteorology/climate science, Chemistry, Electronics/electrical theory, Structural mechanics, Combustion, Information management/statistics, Visualization
  - Most cited sub-disciplines were:
    - Radiation biology, Nuclear physics, Oceanography, Proteomics, Astrophysics, Particle physics, Quantum chemistry, Computer science, Statistics

Related Observations

1. Demand for multidisciplinary simulations is high
2. HPC architecture shifts ⇧ burden on s/w developers
3. Major challenge in application tool development
4. **Major challenge for risk-averse management**
5. **Significant gaps in skillsets and culture**
6. Small Δ in HPC can generate 3-digit ROI
7. **Q: 64-core workstations – are they the answer?**

Industry Engagement is:

DARPA –	applied military technology
DOE –	tech transfer
DoD –	procurement
Commerce –	export support
NSF –	scientific discovery
NCSA –	intersection of HPC and x,y,z

iForge: purpose-built for engineering MS&A

Competency Pillars of HPC-Infused Innovation

Petascale, ISV code, data management, architecture, etc.

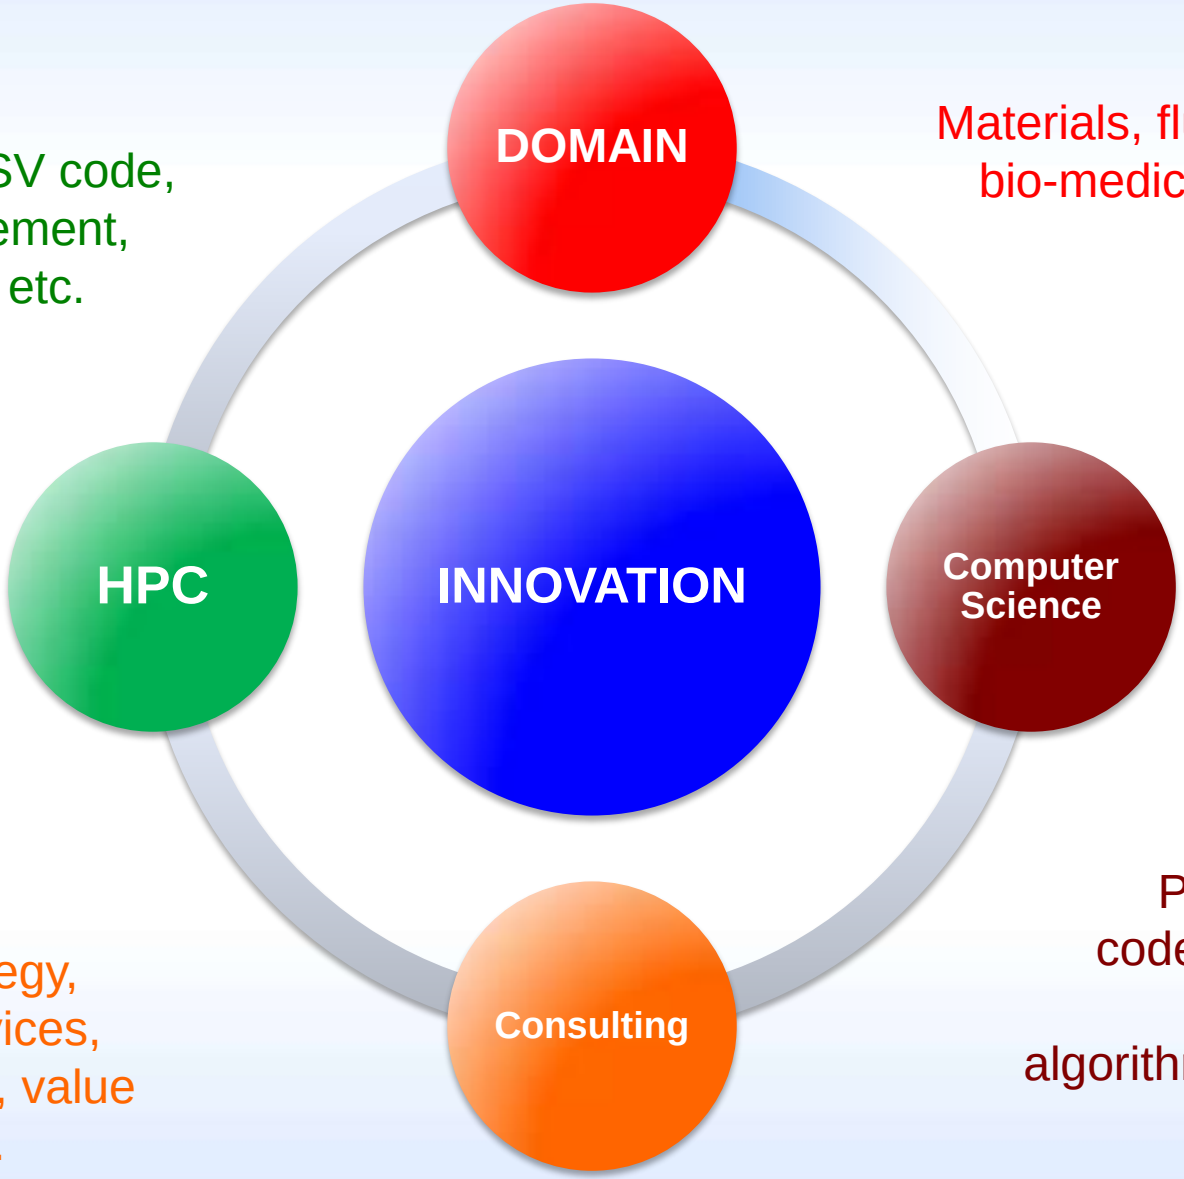

Materials, fluids, thermal, bio-medical, genomics, MRI, etc.

Shared strategy, bilateral services, partnerships, value creation, etc.

Parallelization, code bottlenecks, multiphysics, algorithms, data, etc.

Future Realities (David Houle's Shift Age)

- **Disintermediation**
- Power to the People
- High Tech/**High Touch** – Work/**Place**
- **Human Technology Interface**
- Smart Everything
- **IP** is the Wealth of the Shift Age (??)
- Impact of **Millennials** and Digital Natives
- Data Mining – The Era of **Big Data**
- Vision is paramount
- Be a Morph Corp
- Disruption/Danger/Opportunity

Addressing the Future

- **Computational shift from academic to enterprise**
- SME-OEM-SME digital supply chain will push cloud/HPC
- Leverage manufacturing lessons in other sectors
- **Disintermediation alters the access model to HPC**
- Bifurcation of loosely and tightly integrated systems
- High demand for seamless access to/from cloud
- **No single entity (corp., org., govt. agency) has $$$$**
- Joint ventures and PPPs will be increasingly important
- $2+ Billion needed for next-generation mfg. software
- User interfaces will rule
- **Integrated concurrent manufacturing is revolutionary!**

Thank you!