Microsoft

EXTREME COMPUTING GROUP

Defining the future.

Enabling The Fourth Paradigm

Dan Reed
Corporate Vice President
Technology Strategy and Policy &
Extreme Computing Group

www.hpcdan.org
reed@microsoft.com

Discovery and Innovation 2020

"In the last two decades advances in computing technology, from processing speed to network capacity and the Internet, have revolutionized the way scientists work.

From sequencing genomes to monitoring the Earth's climate, many recent scientific advances would not have been possible without a parallel increase in computing power - and with revolutionary technologies such as the quantum computer edging towards reality, what will the relationship between computing and science bring us over the next 15 years?"

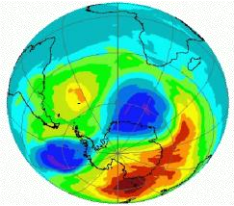

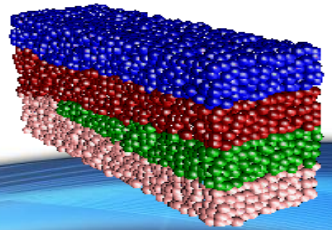

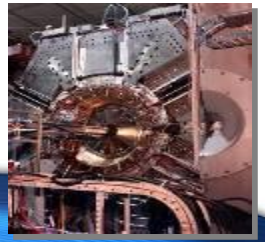

http://research.microsoft.com/towards2020science

Microsoft

The Good News (So Far) …

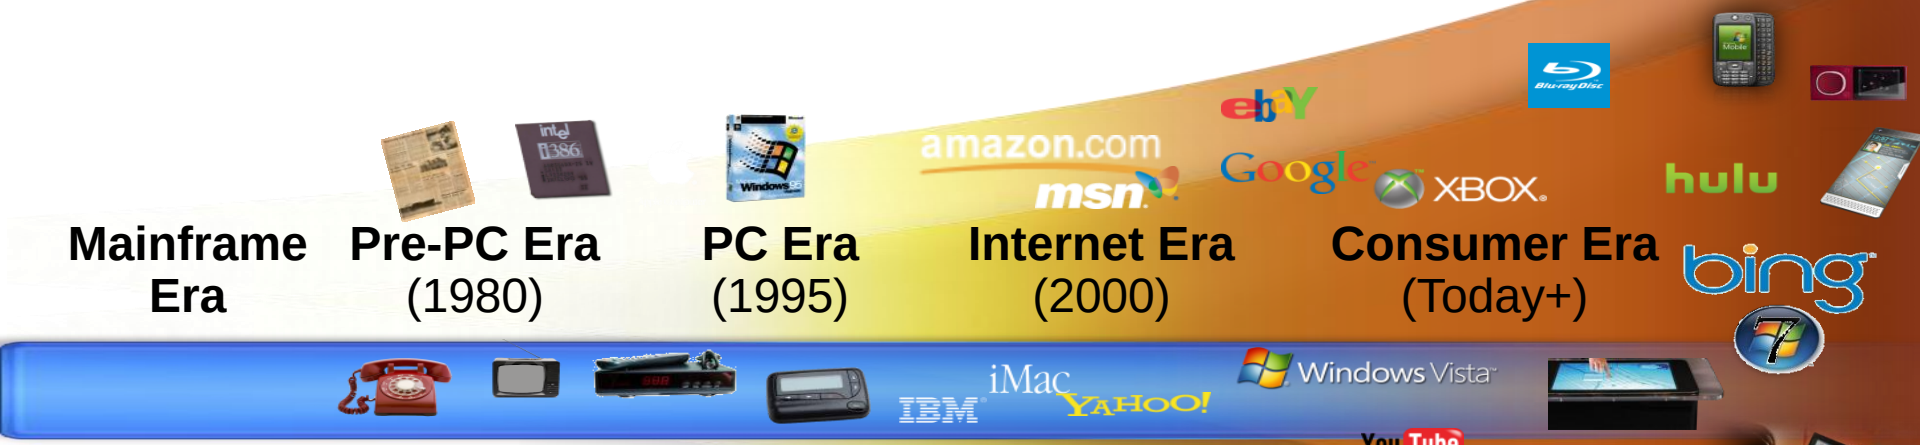

- 21st century implicit computing
  - Increasingly natural interfaces
  - Embedded intelligence
  - Massive data correlation

Microsoft

Convergent Inflections

- Clock rate/power limitations
  - Rise of manycore processors
  - Limited software acceleration via technology alone
  - Clients, servers and infrastructure
- "Surrounded by opportunities"
  - Devices and architectures
  - Programming models and abstractions
  - Algorithms and applications
- From challenge comes opportunity …
  - New applications and systems will arise
    - Including consumer and research

The Data Explosion: The "Other" Exponential

Experiments

Simulations

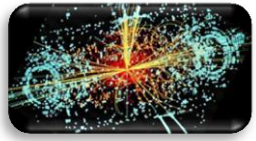

Archives

Literature

Consumer

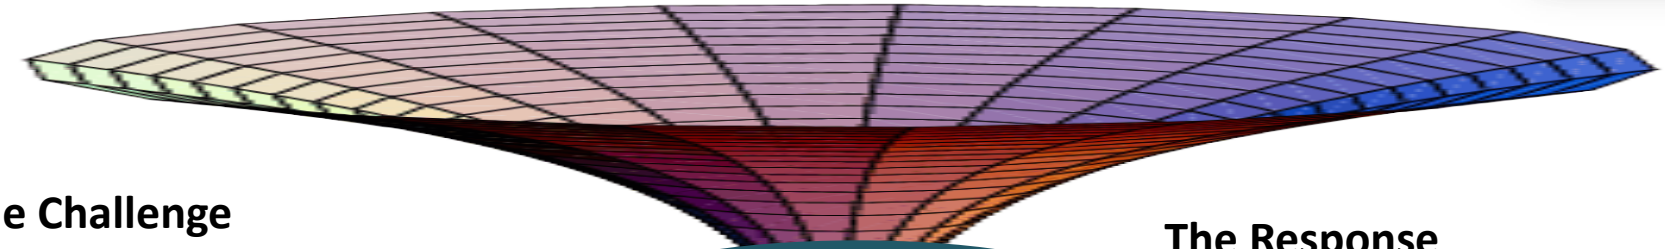

The Challenge

Enable Discovery

Deliver the capability to mine, search and analyze this data in near real time

Petabytes
Doubling & doubling

The Response

A massive private sector build-out of data centers

The Changing Nature of Research

- Thousand years ago – Experimental Science
  - Description of natural phenomena
- Last few hundred years – Theoretical Science
  - Newton's laws, Maxwell's equations…
- Last few decades – Computational Science
  - Simulation of complex phenomena
- Today – Data-centric Science
  - Unify theory, experiment and simulation
  - Using data exploration and data mining
    - Data captured by instruments
    - Data generated by simulations
    - Data generated by sensor networks
    - Data generated by humans

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{4\pi G\rho}{3} - \mathrm{K}\frac{c^2}{a^2}$$

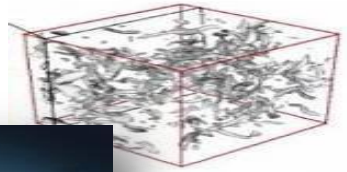

From Genetics To P4 Medicine

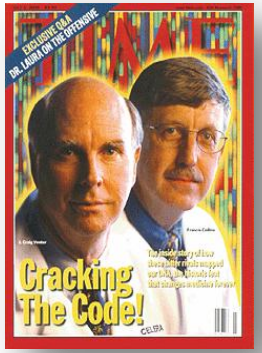

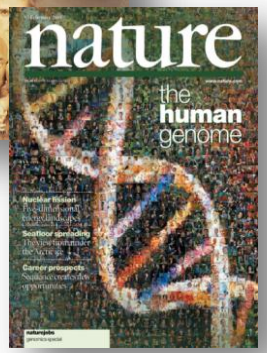

Airway/flow

Mucus

Cilia

Cell biochemistry and structure

Proteomics

Genomics

Lessons From Astronomy

- Historically, discoveries accrued to those
  - With access to unique data
  - Who built next generation telescopes
- Two things changed
  - Growing costs and complexity of telescopes
  - Emergence of whole sky surveys
- The result – virtual astronomy
  - Discovering significant patterns
    - Analysis of rich image/catalog databases
  - Understanding complex astrophysical systems
    - Integrated data/large numerical simulations

Social Implications of the Data Deluge

- Hypothesis-driven
  - "I have an idea, let me verify it."
- Exploratory
  - "What correlations can I glean?"

- Different tools and techniques
  - Rapid exploration of alternatives
  - Data volume and complexity are assets
    - … and challenges

The Fourth Paradigm

- Complex models
  - Multidisciplinary interactions
  - Wide temporal and spatial scales
- Large multidisciplinary data
  - Real-time steams
  - Structured and unstructured
- Distributed communities
  - Virtual organizations
  - Socialization and management
- Diverse expectations
  - Client-centric and infrastructure-centric

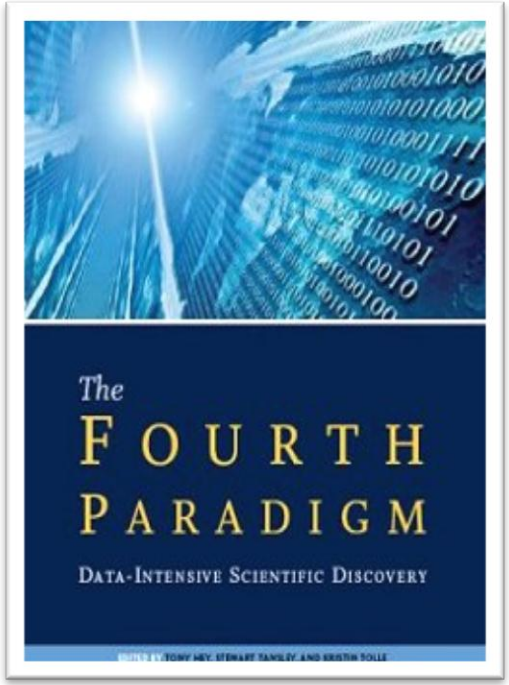

http://research.microsoft.com/en-us/collaboration/fourthparadigm/

Today's Truisms (2009)

- Bulk computing is almost free
  - … but software and power are not
- Inexpensive sensors are ubiquitous
  - … but scientific data fusion remains difficult
- Moving lots of data is {still} hard
  - … because we're missing trans-terabit/second networks
- People are really expensive!
  - … and robust software remains extremely labor intensive
- Innovation challenges are changing
  - … and the technology must empower, not frustrate

Free Storage: Like Free Puppies

- Storage is cheap (<<$1K/TB)
- Storage management is not
  - OPEX > 100 CAPEX
  - Goal: OPEX << CAPEX

Research Empowerment Challenges

- Optimize for human creativity
  - Seamlessly accessible
  - Intentional not imperative
  - Anticipatory not reactive

- Insatiable infrastructure demand
  - Cycles, storage, support
- Distributed acquisition/deployment
  - Duplicative, non-shared infrastructure
- Distributed cost structures
  - Power, space, staff, staff, hardware
- Long-term sustainability
  - Decades rather than months/years

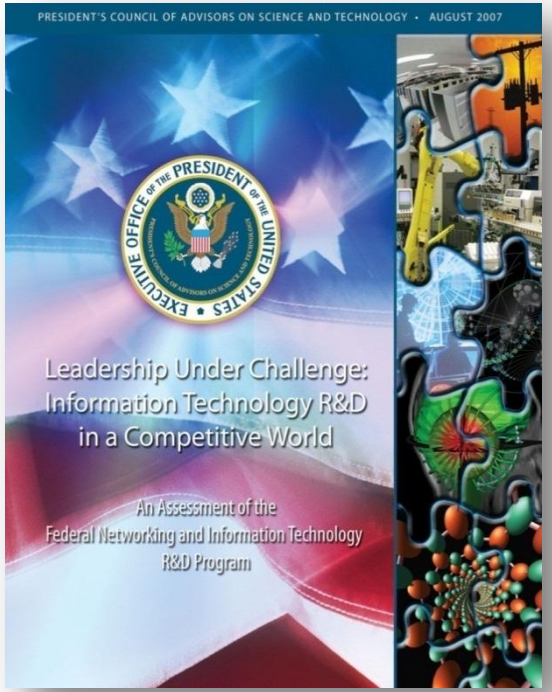

The "Branscomb" Computing Pyramid

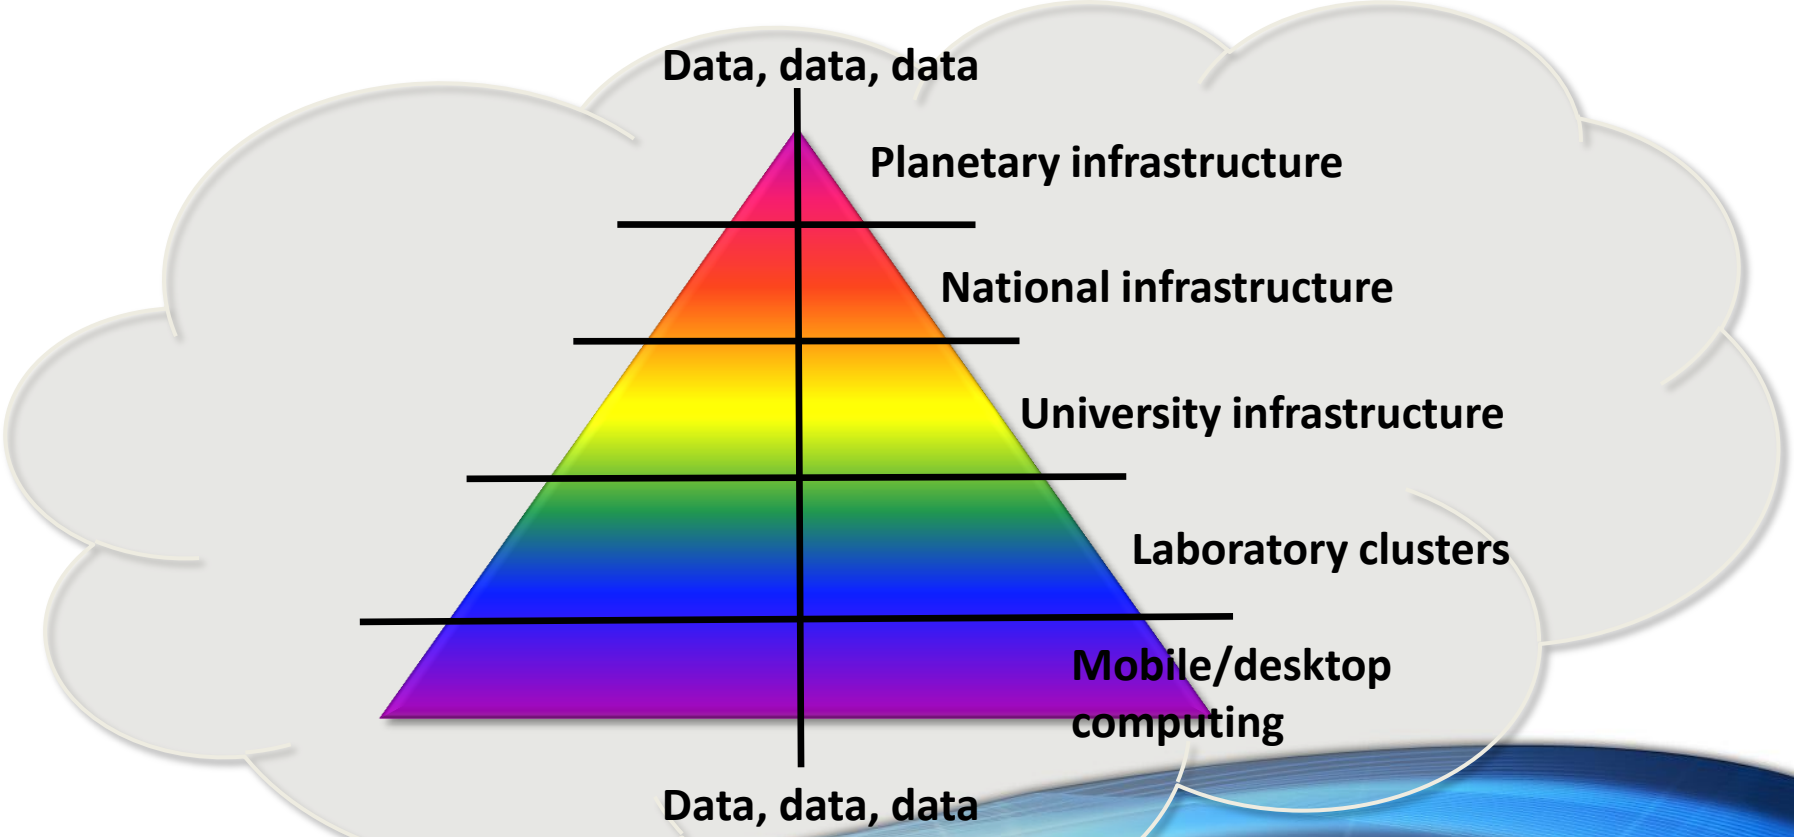

Orders of Magnitude Always Matter

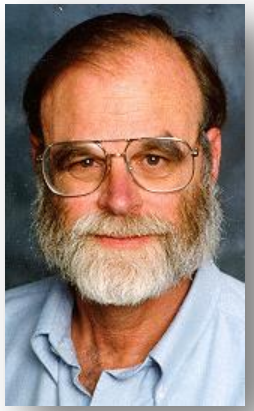

- Multidisciplinary challenges are the present and future
  - … and the tools must empower, not frustrate
- These are systemic problems

- An insight from Jim Gray …

- A computation task has four characteristic demands:
  - Networking – delivering questions and answers
  - Computation – transforming information to produce new information
  - Data access – access to information needed by the computation
  - Data storage – long term storage of information
- The ratios among these *and their costs* are critical

Clouds and Hosted Infrastructure

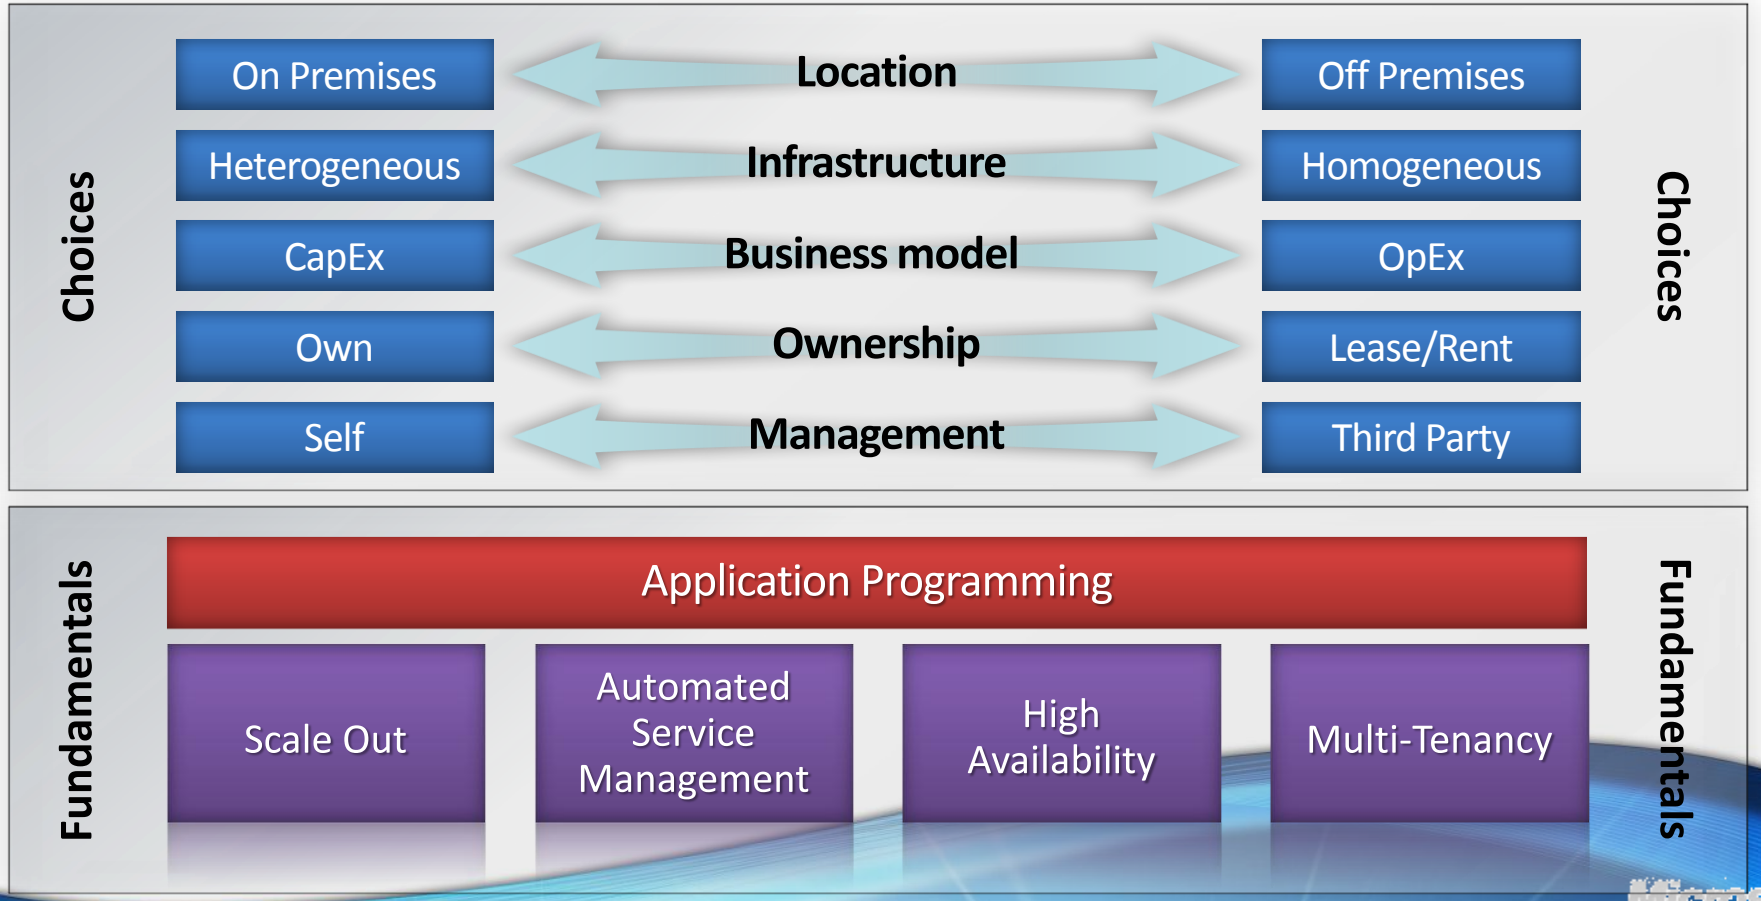

Generic Computer Physical Plant

- Massive commodity servers
- Energy intensive infrastructure
- Cooling inefficiencies
- Environmental issues
- Expensive UPS support
- Enterprise TCP/IP networks
- Long deployment times
- Diverse services and SLAs

- Many optimization opportunities …

Standard IT and Clouds: Twins Separated At Birth

- Similar technology issues but vastly different scales
  - Node and system architectures
  - Communication fabrics
  - Storage systems and analytics
  - Physical plant and operations
  - Programming models
  - Reliability and resilience
- Differing culture and sociology
  - Design and operations
  - Management and philosophy

Cloud Data Centers: Scale Calibration …

Each data center is
11.5 times
the size of a football field
and consumes 40 MW at the
utility meter

A Computer Room Is Not A Data Center

Data Center Facility "PacMan"

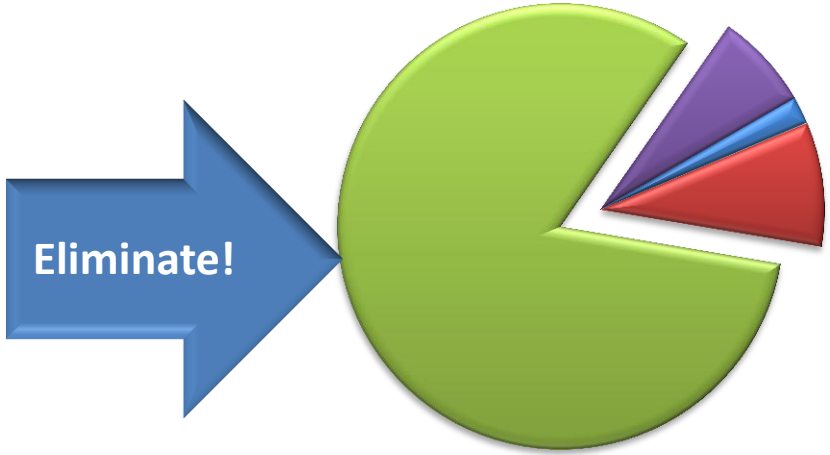

- Land - 2%
- Core and shell costs – 9%
- Architectural – 7%
- Mechanical/Electrical – 82%
  - 16% increase/year since 2004

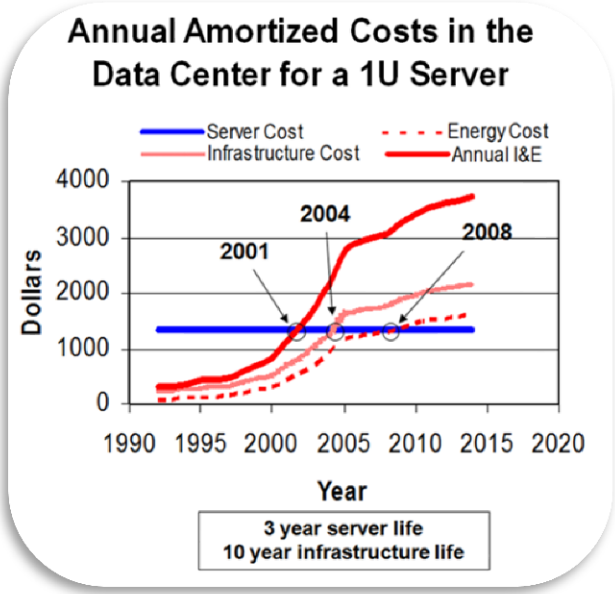

Belady, C., "In the Data Center, Power and Cooling Costs More than IT Equipment it Supports", *Electronics Cooling Magazine* (February 2007)

Source: Christian Belady

Current Facilities Economic Reality …

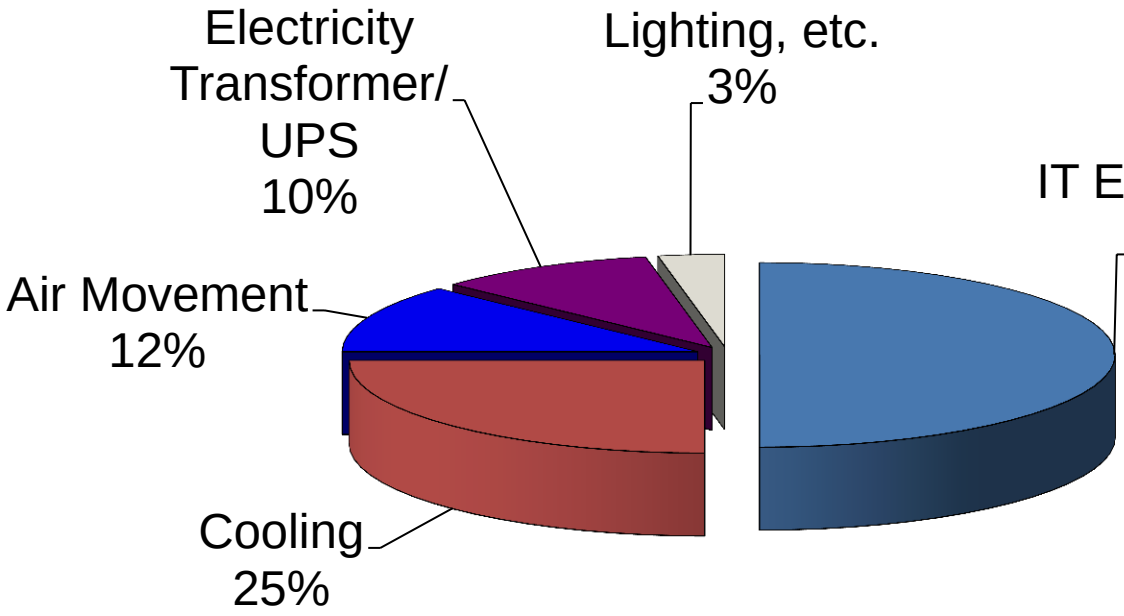

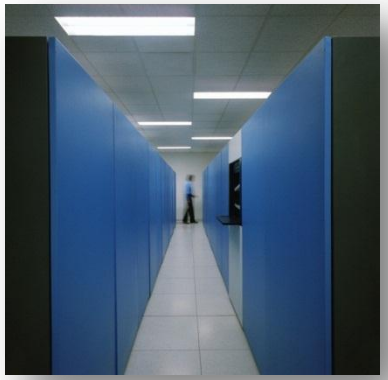

Source: EYP Mission Critical Facilities Inc., New York

Systemic Design: Not Just Watts …

- Watts alone are irrelevant
  - Turn off the equipment and declare victory
- The real metric is the following …

$$\frac{\textit{Effective Operations}}{\textit{Total Cost of Ownership}}$$

- Many convolved ideas
  - Application types and needs
  - Microarchitecture and system design
  - Power distribution efficiency
  - Packaging and cooling overhead
  - Market costs for power and hardware
  - Cost of people and money

Containers and Efficiency

Rethinking Packaging and Cooling

- People and hardware need not mix
- Hardware cooling standards are conservative
  - Reliable at high temperature/humidity
- Optimize for efficiency
  - Cooling is (often) unnecessary
    - Design for ambient environments
  - Energy reliability is (often) unnecessary
    - Design for power outages
  - Use larger building blocks
    - Accept component failures

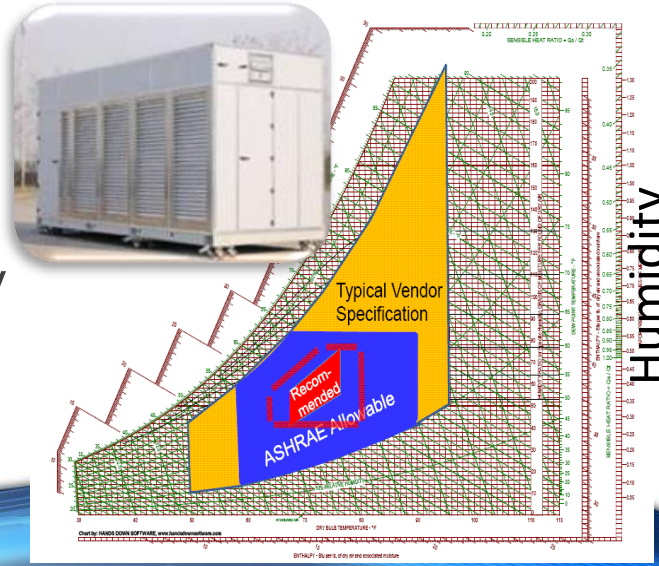

Microsoft

Microsoft Gen4 Data Centers

- Scalable
- Plug-and-play spine infrastructure
- Factory pre-assembled
  - Pre-assembled containers (PACs)
  - Pre-manufactured buildings (PMBs)
- Rapid deployment
- De-mountable
- Reduced construction
- Sustainable measures
- Map applications to class

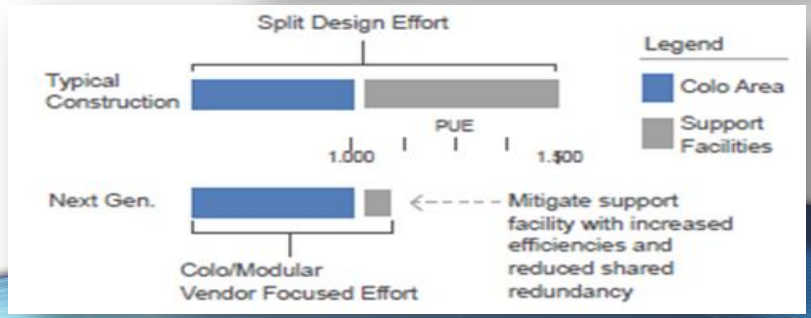

U.S. NITRD Program Evaluations

- PITAC's 1999 overall assessment
  - *Information Technology Research: Investing in Our Future*
- During 2003-2005, focused PITAC assessments
  - Health care and IT, cybersecurity
  - Computational science
- PCAST 2007 review
  - Successor to 1999 assessment

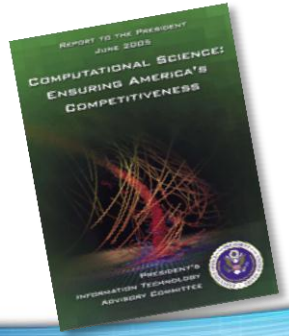

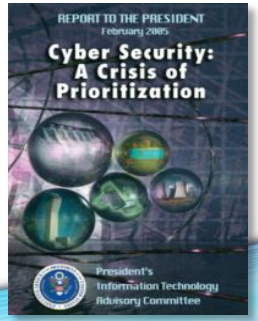

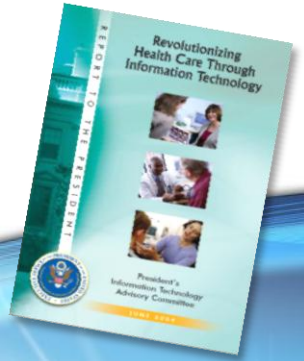

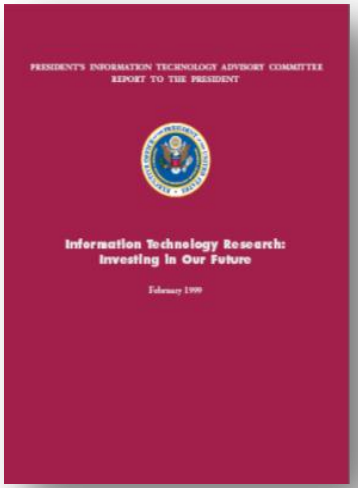

Microsoft

PCAST Principal Findings (2007)

- The United States is the current global leader in networking and information technology
- That leadership is essential to U.S. economic prosperity, security, and quality of life
- It is the product of the entire U.S. NIT ecosystem – industry, government, and academia, with a key role played by Federal R&D support
- But it is being challenged by other nations – not only established competitors in Asia and Western Europe, but also newcomers such as India and China – that are investing to build strong NIT ecosystems
- The nature and scale of Federal NIT R&D coordination processes are inadequate to support continued U.S. leadership

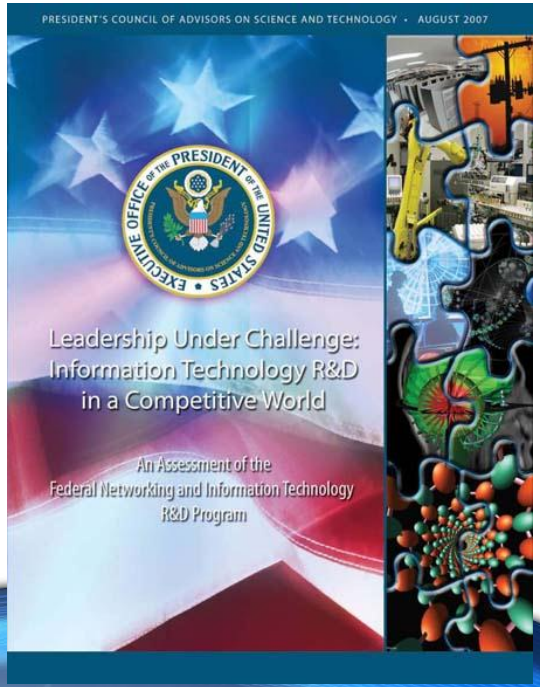

PCAST Principal Recommendations (2007)

- In response to the competitive challenge, the U.S. must:
- Revamp NIT education and training
  - New curricula and approaches to meet demands
  - Increased fellowships and streamlined visa processes
- Rebalance the Federal NIT R&D portfolio
  - More long-term, large-scale, multidisciplinary R&D
  - More innovative, higher-risk R&D
- Reprioritize the Federal NIT R&D topics
  - Increase: systems connected with physical world, software, digital data, and networking
  - Sustain: high-end computing, security, HCI, and social sciences
- Improve planning and coordination of Federal NIT R&D programs

Data Stores and Data Streams: PCAST 2007

- Findings
  - Volume of digital data is an opportunity to advance U.S. leadership in science and technology – harnessing it is a national priority
  - "Data deluge" is overwhelming the capacity of academic institutions and Federal agencies. More robust NIT capabilities are needed
- Recommendations
  - The Interagency Working Group on Digital Data, with the NITRD Subcommittee, should develop a national strategy and develop and implement a plan to assure the long-term preservation, stewardship, and widespread availability of data important to S&T
  - As part of this effort, NITRD program agencies should develop a multi-agency plan for coordinated R&D to advance data management and analysis

Fourth Paradigm: Seeding Change

- New public and private sector partnerships
  - Government, academia, NGOs, industry …
- Core competency foci
  - Innovation and infrastructure
- Sustainability and relationships
  - Prototyping versus long-term partnerships
- Multidisciplinary scaling and fusion
  - Co-location versus distribution
- Security, privacy and provenance
  - Insights versus exposure

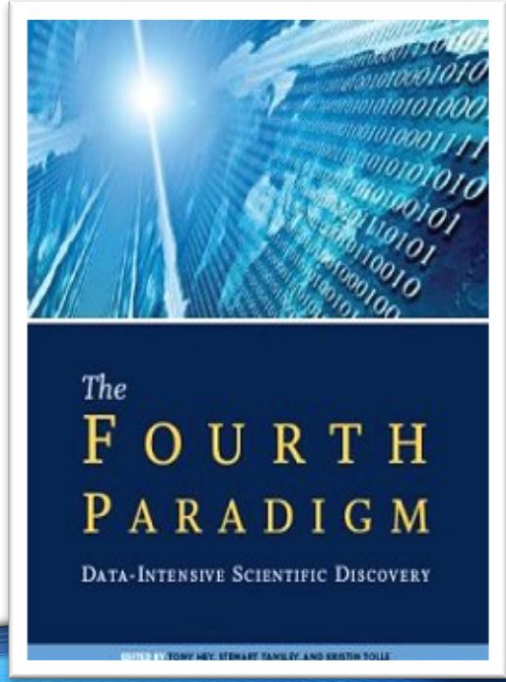

Today Is An Inflection Point

- Economic challenges
  - Research efficiency
  - Infrastructure sustainability
- Innovation opportunities
  - Multidisciplinary data fusion
  - Deep data mining
- Technology transition
  - Scaling economics
  - Rich cloud and web services

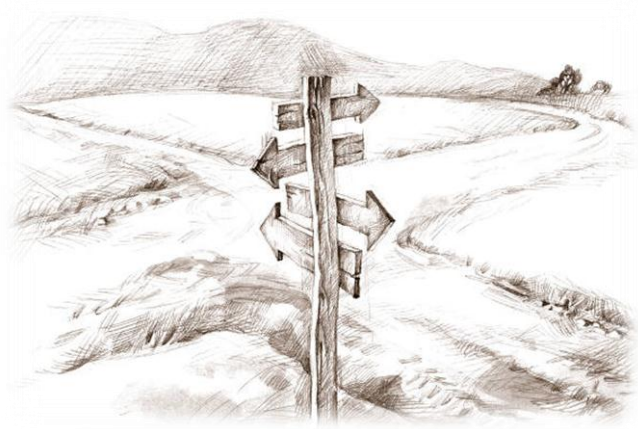

Microsoft®

Your potential. Our passion.™