

IBM Blue Gene/P

Dr. George Chiu
IEEE Fellow
IBM T.J. Watson Research Center
Yorktown Heights, NY



President Obama Honors IBM's Blue Gene Supercomputer With National Medal Of Technology And Innovation

*Ninth time IBM has received nation's most prestigious tech award
Blue Gene has led to breakthroughs in science, energy efficiency and analytics*



WASHINGTON, D.C. - 18 Sep 2009: President Obama recognized IBM (NYSE: [IBM](#)) and its *Blue Gene* family of supercomputers with the National Medal of Technology and Innovation, the country's most prestigious award given to leading innovators for technological achievement. President Obama will personally bestow the award at a special White House ceremony on October 7. IBM, which earned the National Medal of Technology and Innovation on eight other occasions, is the only company recognized with the award this year.

Blue Gene's speed and expandability have enabled business and science to address a wide range of complex problems and make more informed decisions -- not just in the life sciences, but also in astronomy, climate, simulations, modeling and many other areas. Blue Gene systems have helped map the human genome, investigated medical therapies, safeguarded nuclear arsenals, simulated radioactive decay, replicated brain power, flown airplanes, pinpointed tumors, predicted climate trends, and identified fossil fuels – all without the time and money that would have been required to physically complete these tasks.

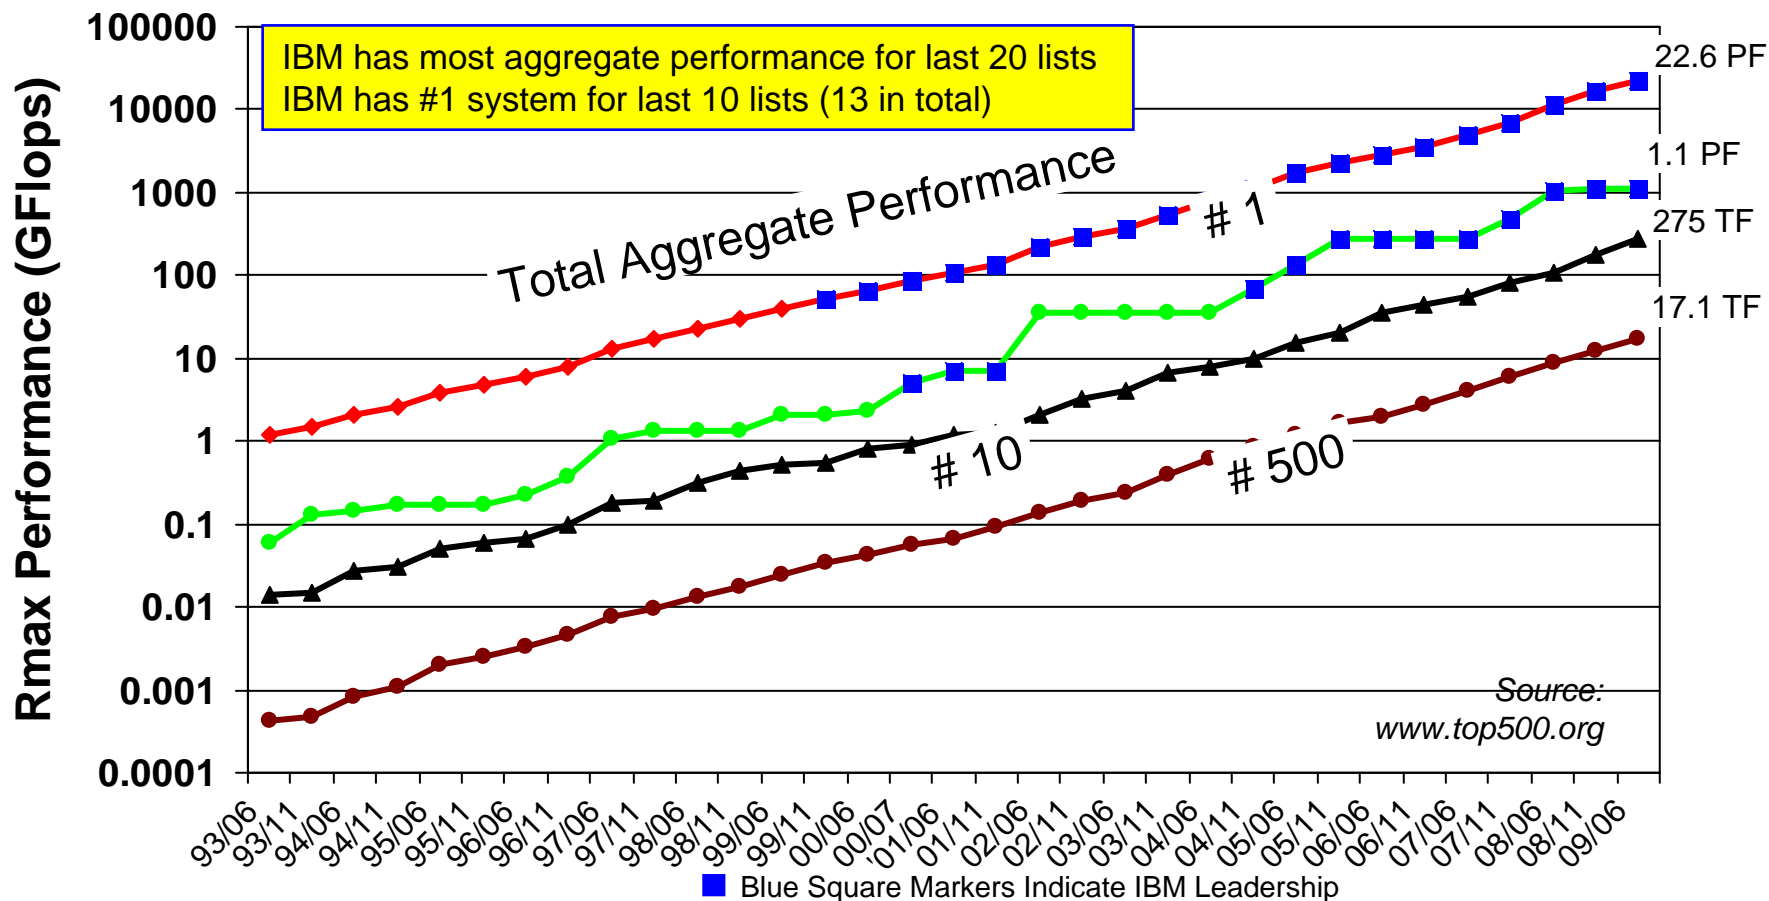
The system also reflects breakthroughs in energy efficiency. With the creation of Blue Gene, IBM dramatically shrank the physical size and energy needs of a computing system whose processing speed would have required a dedicated power plant capable of generating power to thousands of homes.

The influence of the Blue Gene supercomputer's energy-efficient design and computing model can be seen today across the Information Technology industry. Today, 18 of the top 20 most energy efficient supercomputers in the world are built on IBM high performance computing technology, according to the latest Supercomputing 'Green500 List' announced by [Green500.org](#) in July, 2009.

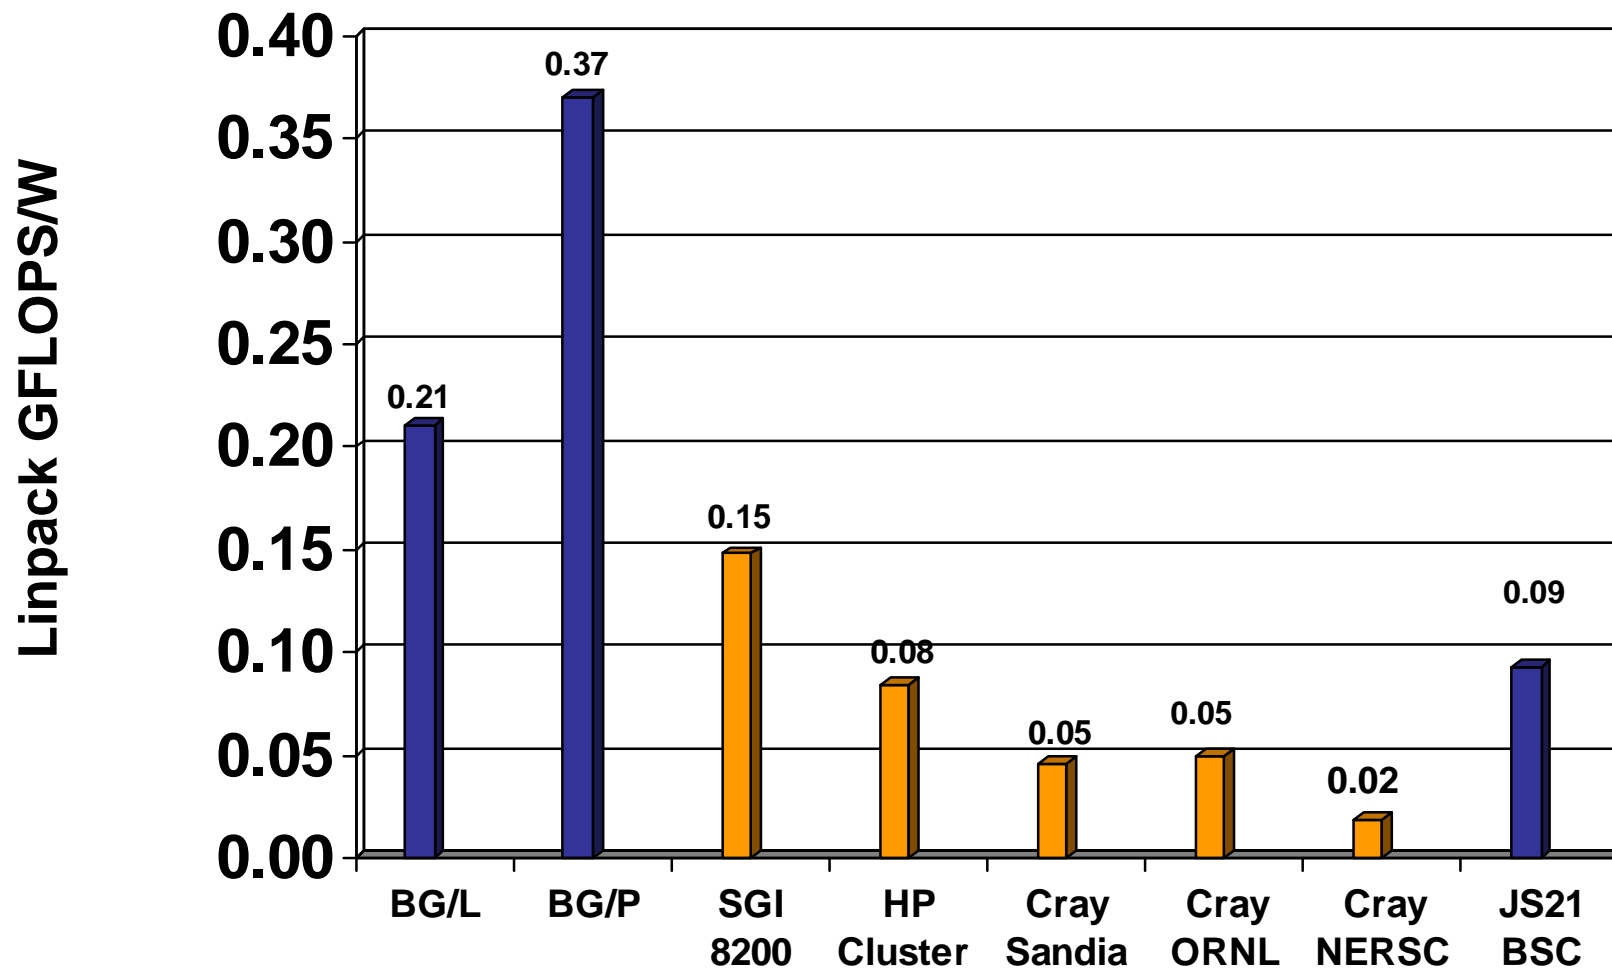
BlueGene Roadmap

- **BG/L (5.7 TF/rack) – 130nm ASIC (1999-2004GA)**
 - 104 racks, 212,992 cores, 596 TF/s, 210 MF/W; dual-core system-on-chip,
 - 0.5/1 GB/node
- **BG/P (13.9 TF/rack) – 90nm ASIC (2004-2007GA)**
 - 72 racks, 294,912 cores, 1 PF/s, 357 MF/W; quad core SOC, DMA
 - 2/4 GB/node
 - SMP support, OpenMP, MPI
- **BG/Q**
 - 20 PF/s

TOP500 Performance Trend



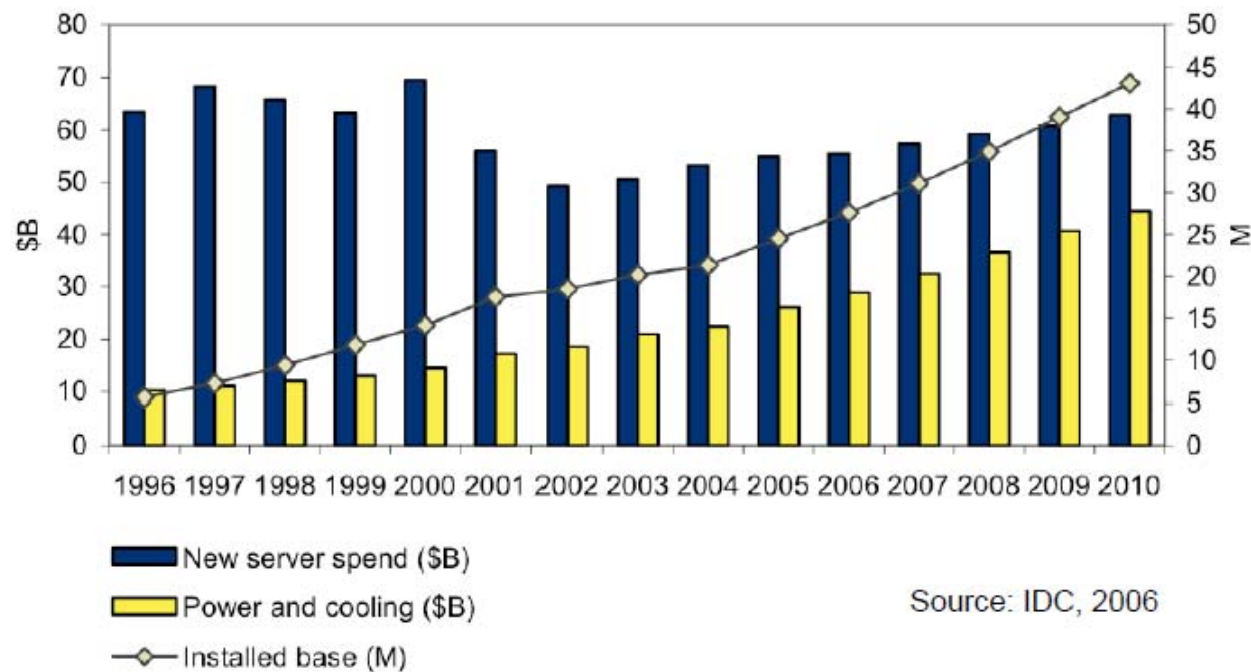
November 2007 Green 500



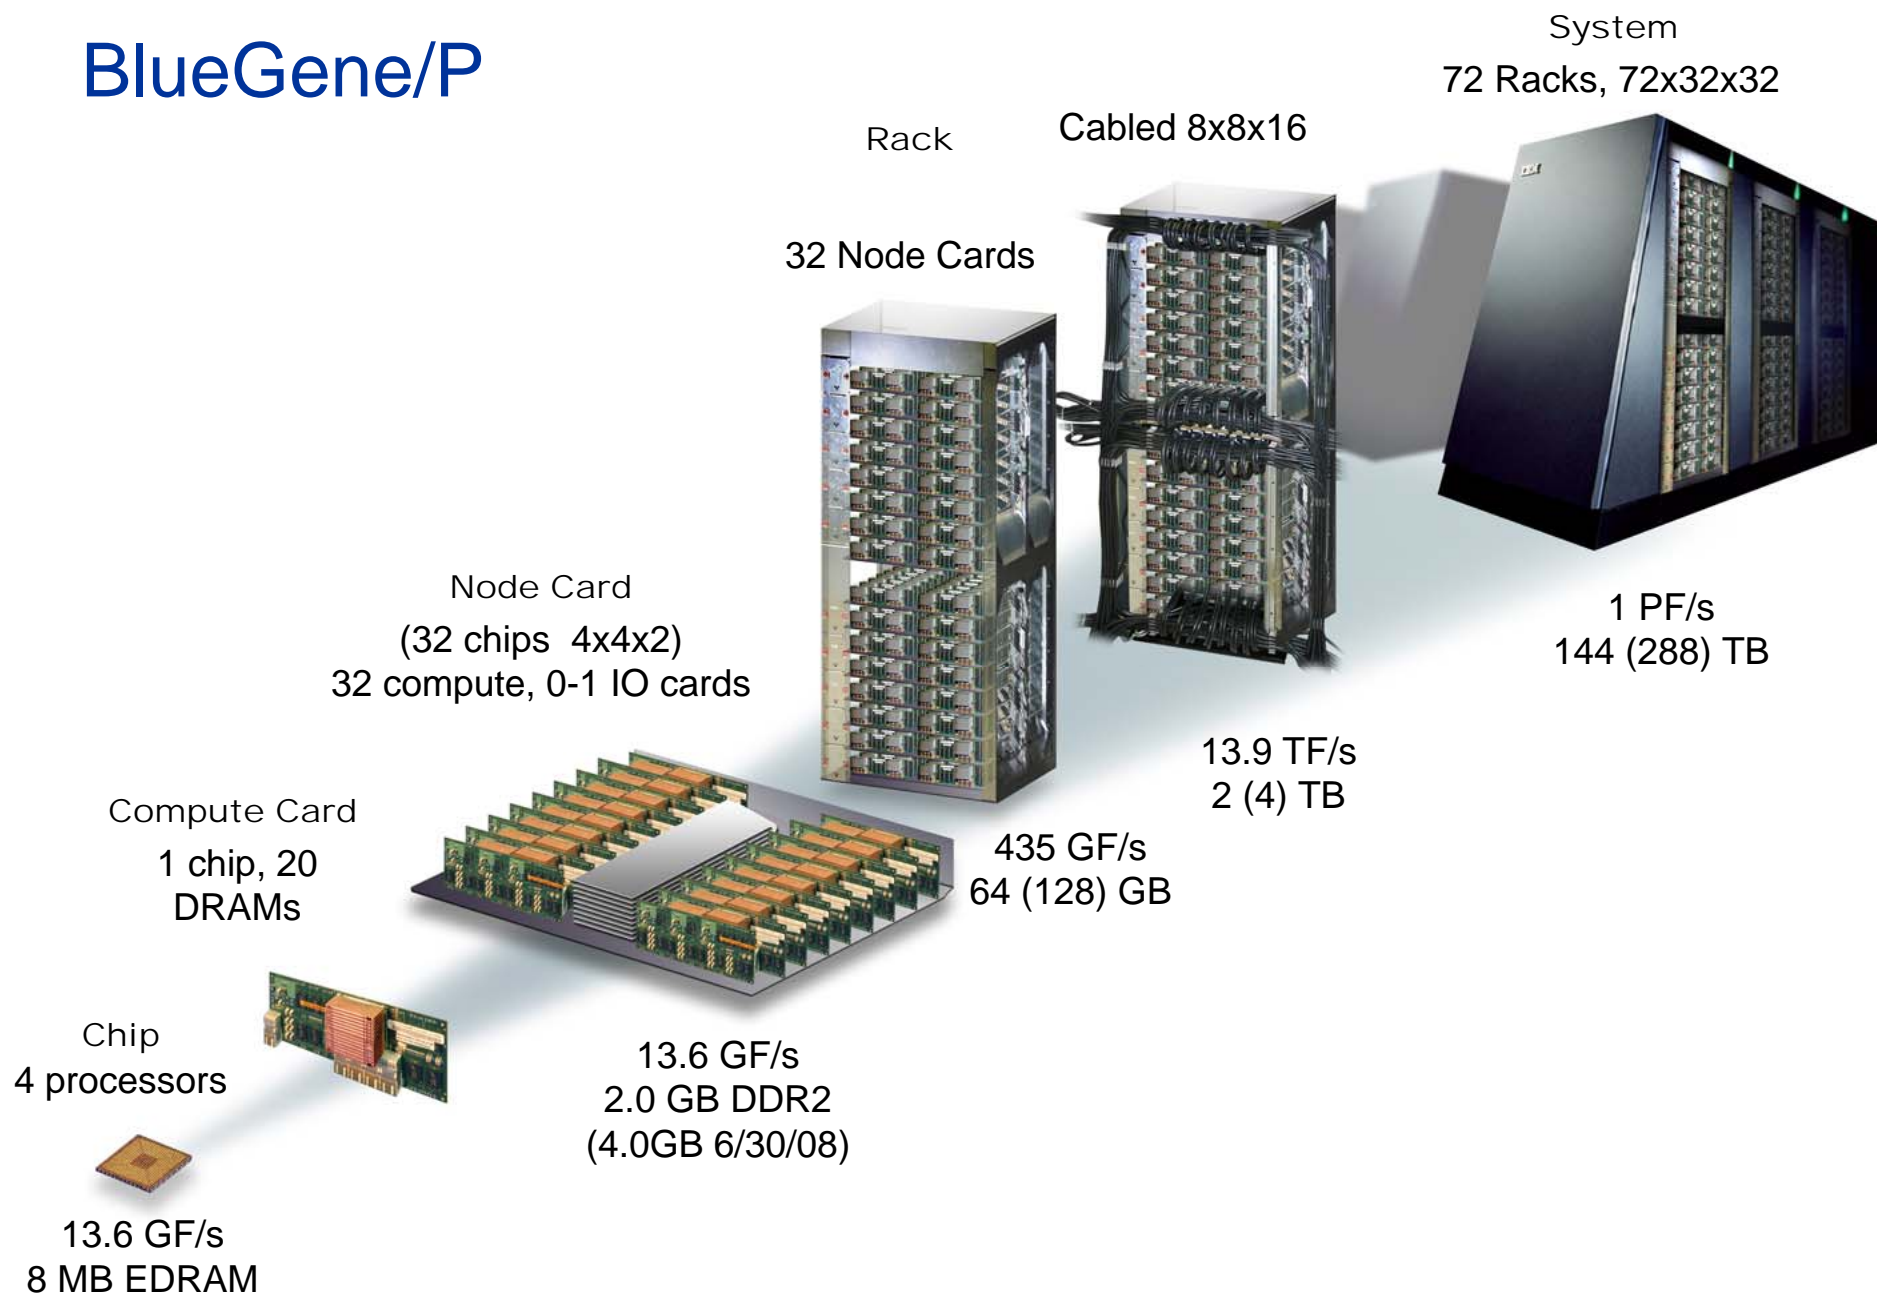
The “Cost” of Achieving Performance

- **Electrical power costs \$\$\$\$.**

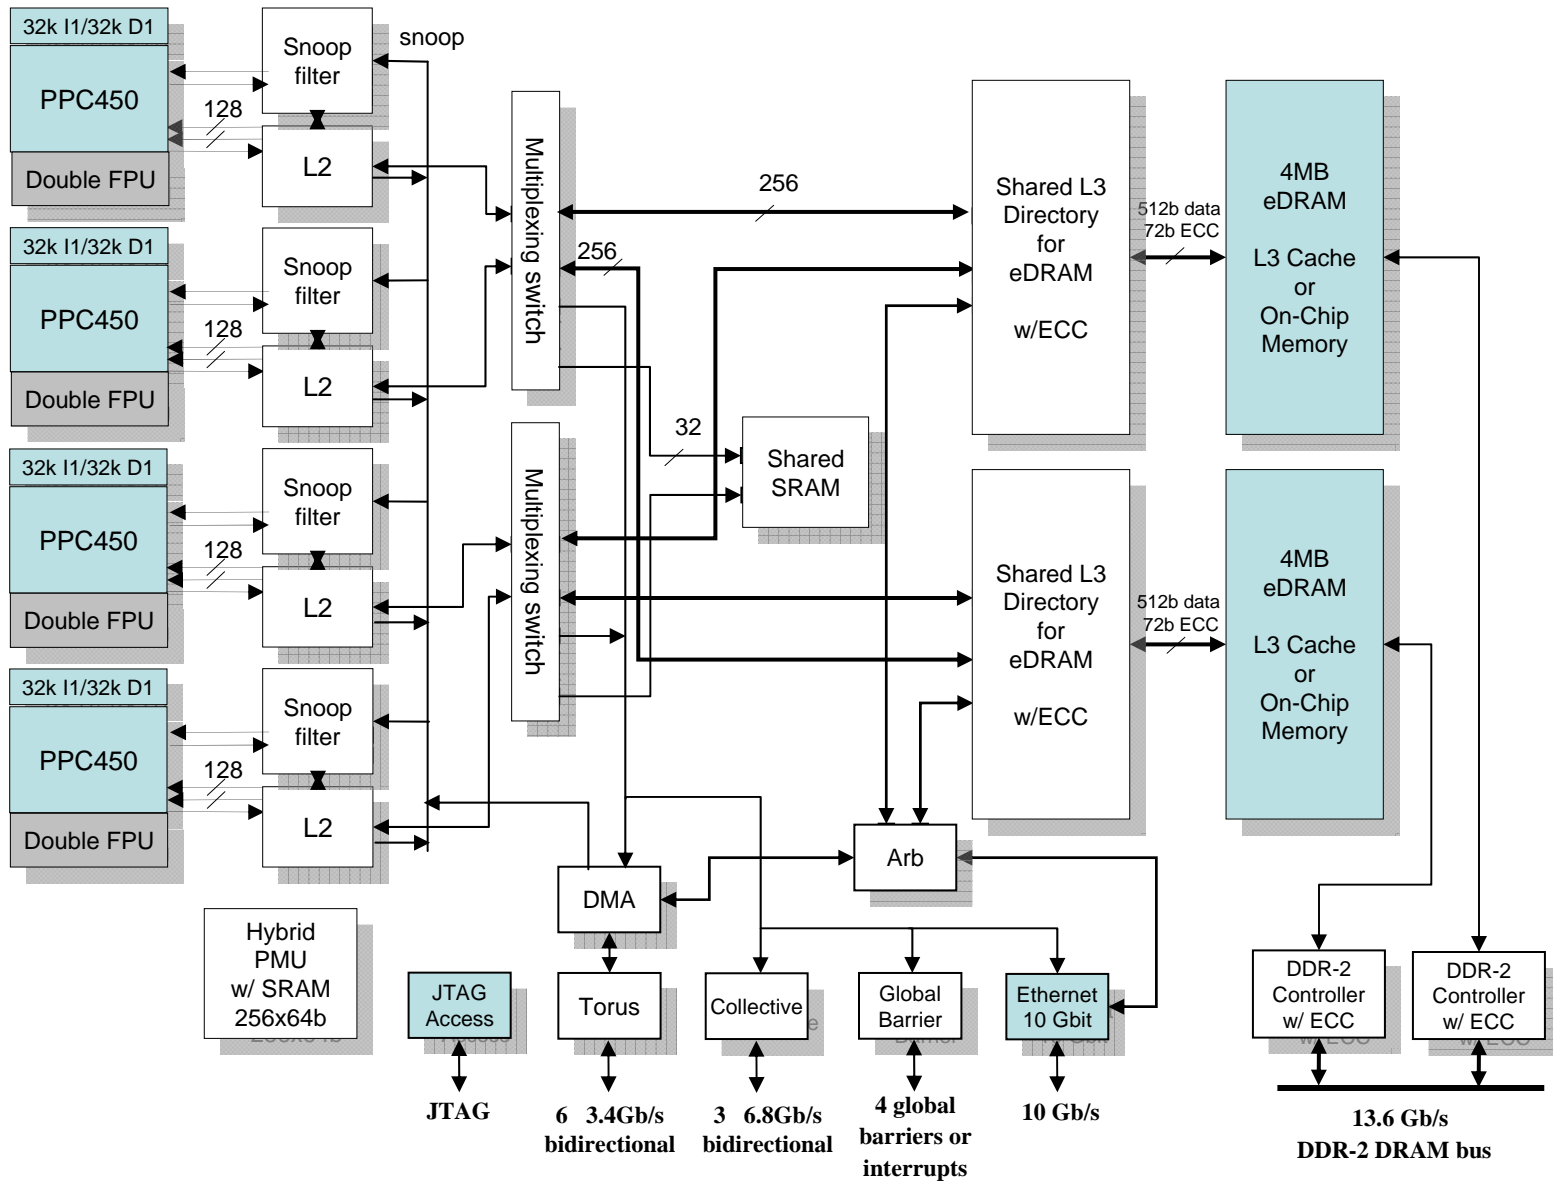
Worldwide Expense to Power and Cool the Server Installed
Base, 1996–2010



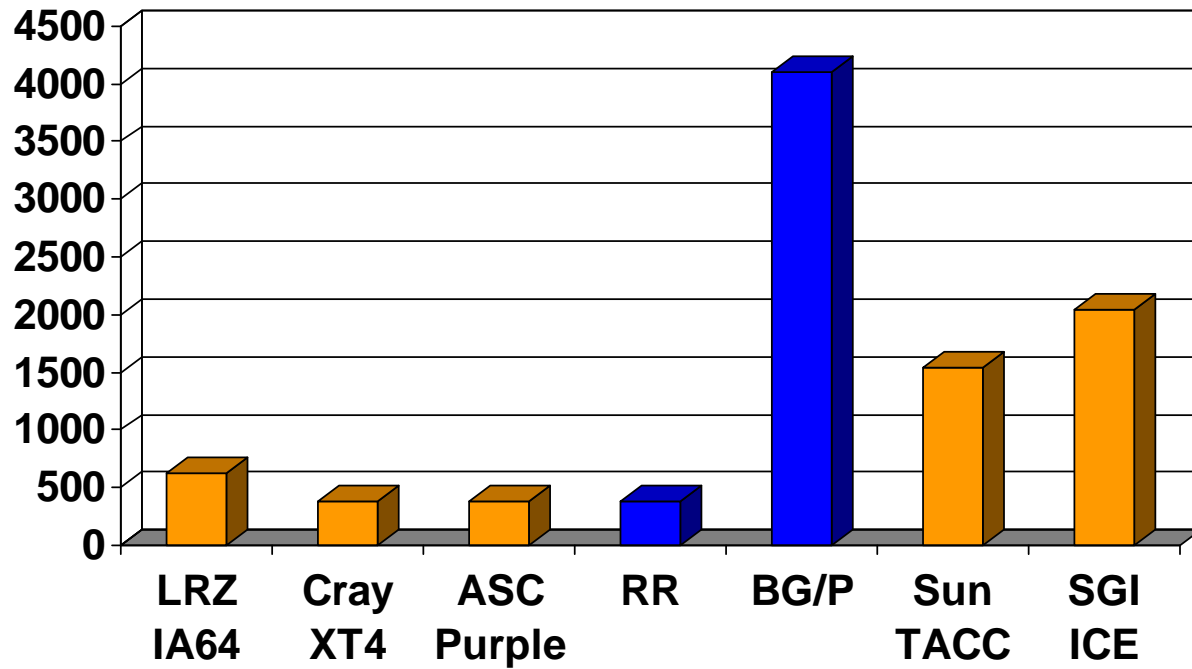
BlueGene/P



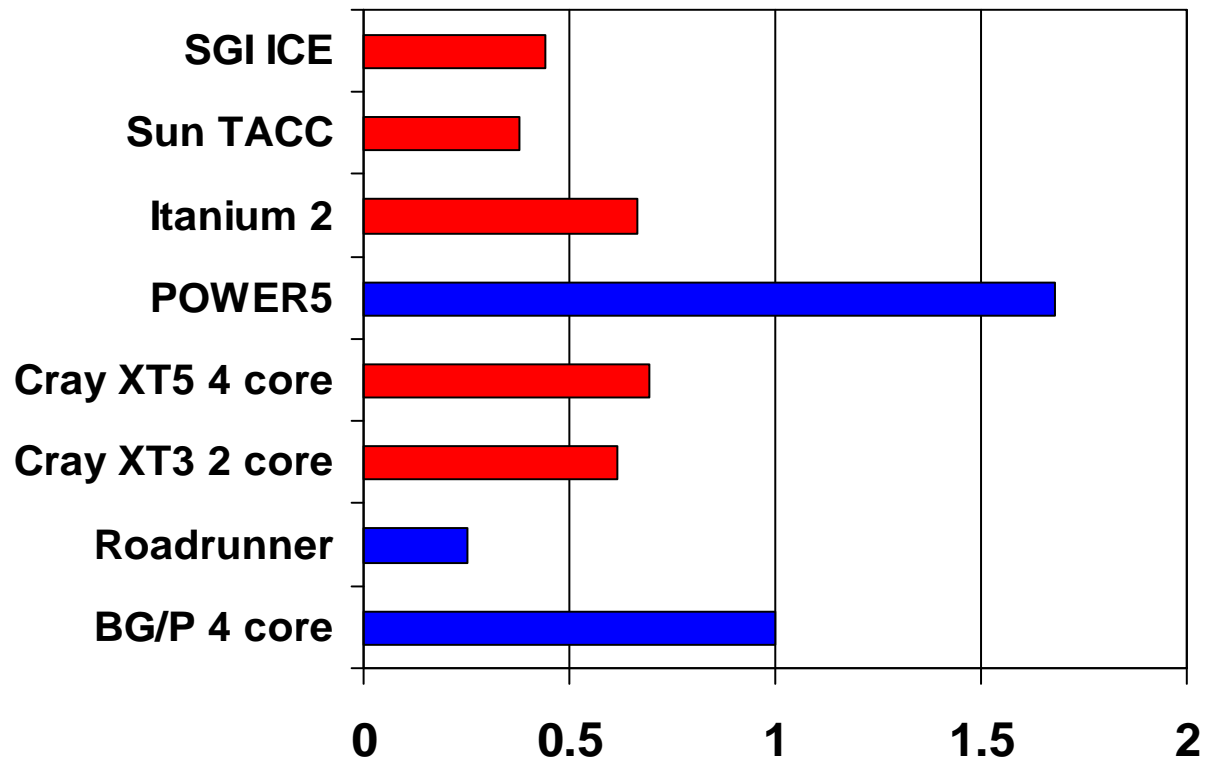
BlueGene/P compute ASIC




Main Memory Capacity per Rack




Peak Memory Bandwidth per node (byte/flop)




Main Memory Bandwidth per Rack



Interprocessor Peak Bandwidth per node (byte/flop)



Failures per Month per @ 100 TFlops (20 BG/L racks) unparalleled reliability



Results of survey conducted by Argonne National Lab on 10 clusters ranging from 1.2 to 365 TFlops (peak); excluding storage subsystem, management nodes, SAN network equipment, software outages

Reproducibility of Floating Point Operations

- The example below is illustrated with single precision floating point operations (~7 digits accuracy), the same principle applies to double precision floating point operations (~14 digits accuracy)
- $A = 1234567$
- $B = 1234566$
- $C = 0.1234567$
- $A - B + C = 1.123457$
- $A + C - B = 1$
- **Caution: floating point with finite number of digits in accuracy is not commutative, and not associative.**
- **BG/L,P enforces execution order, so all calculations are reproducible.**

Summary

- **Blue Gene/P: Facilitating Extreme Scalability**
 - Ultrascale capability computing when nothing else will satisfy
 - Provides customer with enough computing resources to help solve grand challenge problems
 - Provide competitive advantages for customers' applications looking for extreme computing power
 - Energy conscious solution supporting green initiatives
 - Familiar open/standards operating environment
 - Simple porting of parallel codes

- **Key Solution Highlights**
 - Leadership performance, space saving design, low power requirements, high reliability, and easy manageability