Petascale Delivered – What's Past is Prologue

IBM's *pNext*; The Next Era of Computing Innovation

A mix of market forces, technological innovations and business conditions have created a tipping point in HPC

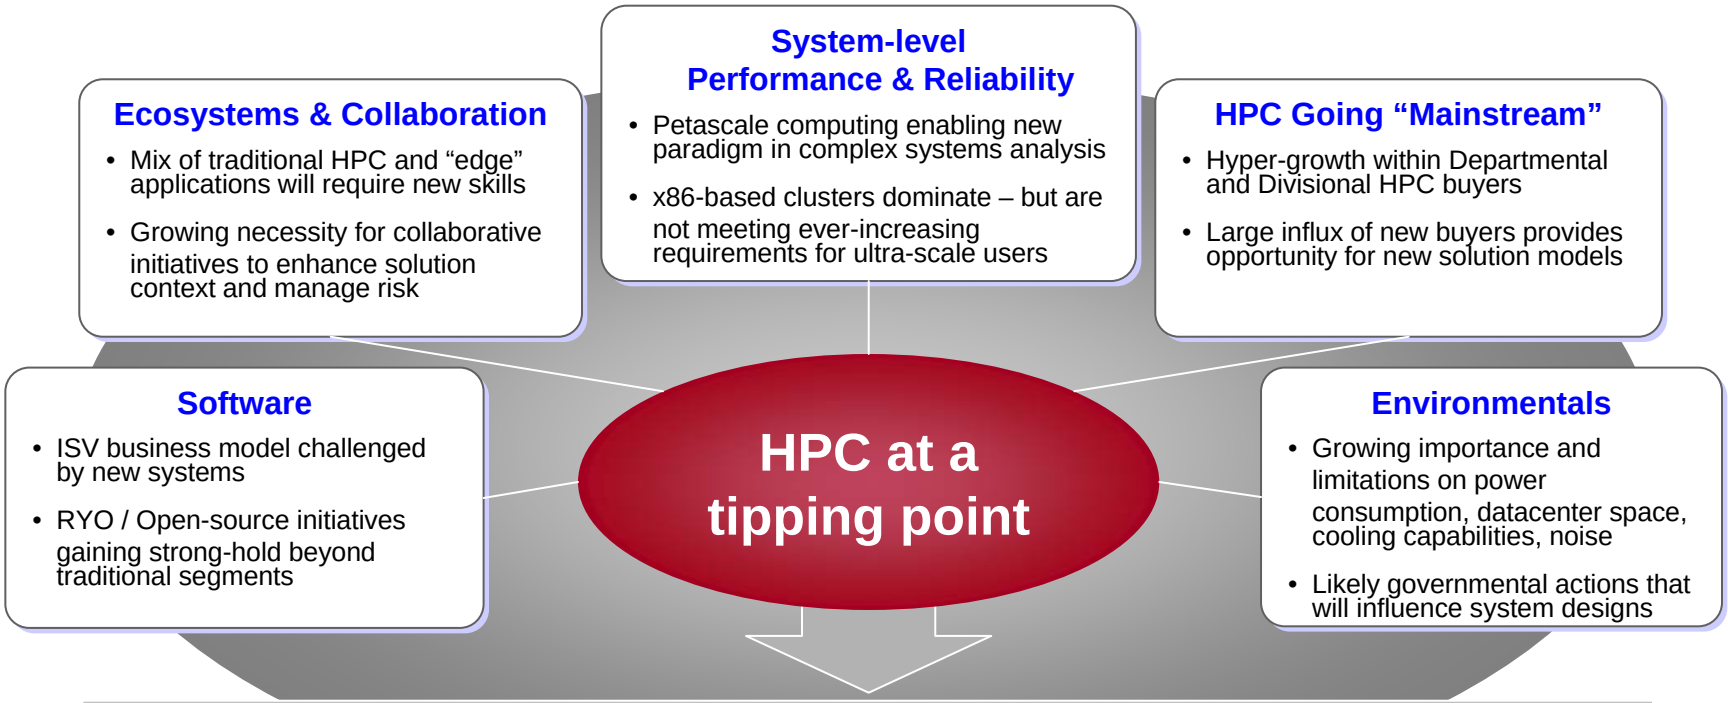

How can we exploit these shifts to meet our Performance & Productivity requirements?

Performance and Productivity Challenges require a Multi-Dimensional Approach

Highly Productive Systems

POWER

HPCS PERCS

Highly Scalable Multi-core Systems

BLUE GENE

Hybrid Systems

Comprehensive (Holistic) System Innovation & Optimization

June 18, 2008 – Petascale Delivered!

Roadrunner Breaks Petaflop Milestone

Date added: 18 Jun 2008

Lead engineer Don Grice of IBM inspects the world's fastest computer in the company's Poughkeepsie, NY plant. The computer nicknamed "Roadrunner" was built for the Department of Energy's National Nuclear Security Administration and will be housed at Los Alamos National Laboratory in New Mexico. IBM engineers in Poughkeepsie, N.Y., Rochester, Minn., Austin, Texas and Yorktown Heights, N.Y., worked on the computer, the first to break a milestone known as a "petaflop" -- the ability to calculate 1,000-trillion operations every second. The computer packs the power of 100,000 laptops -- a stack 1.5 miles high. Roadrunner will primarily be used to ensure national security, but will also help scientists perform research into energy, astronomy, genetics and climate change.

Groundbreaking system will have a profound impact on science, business and society

The making of Roadrunner – Key Building Blocks

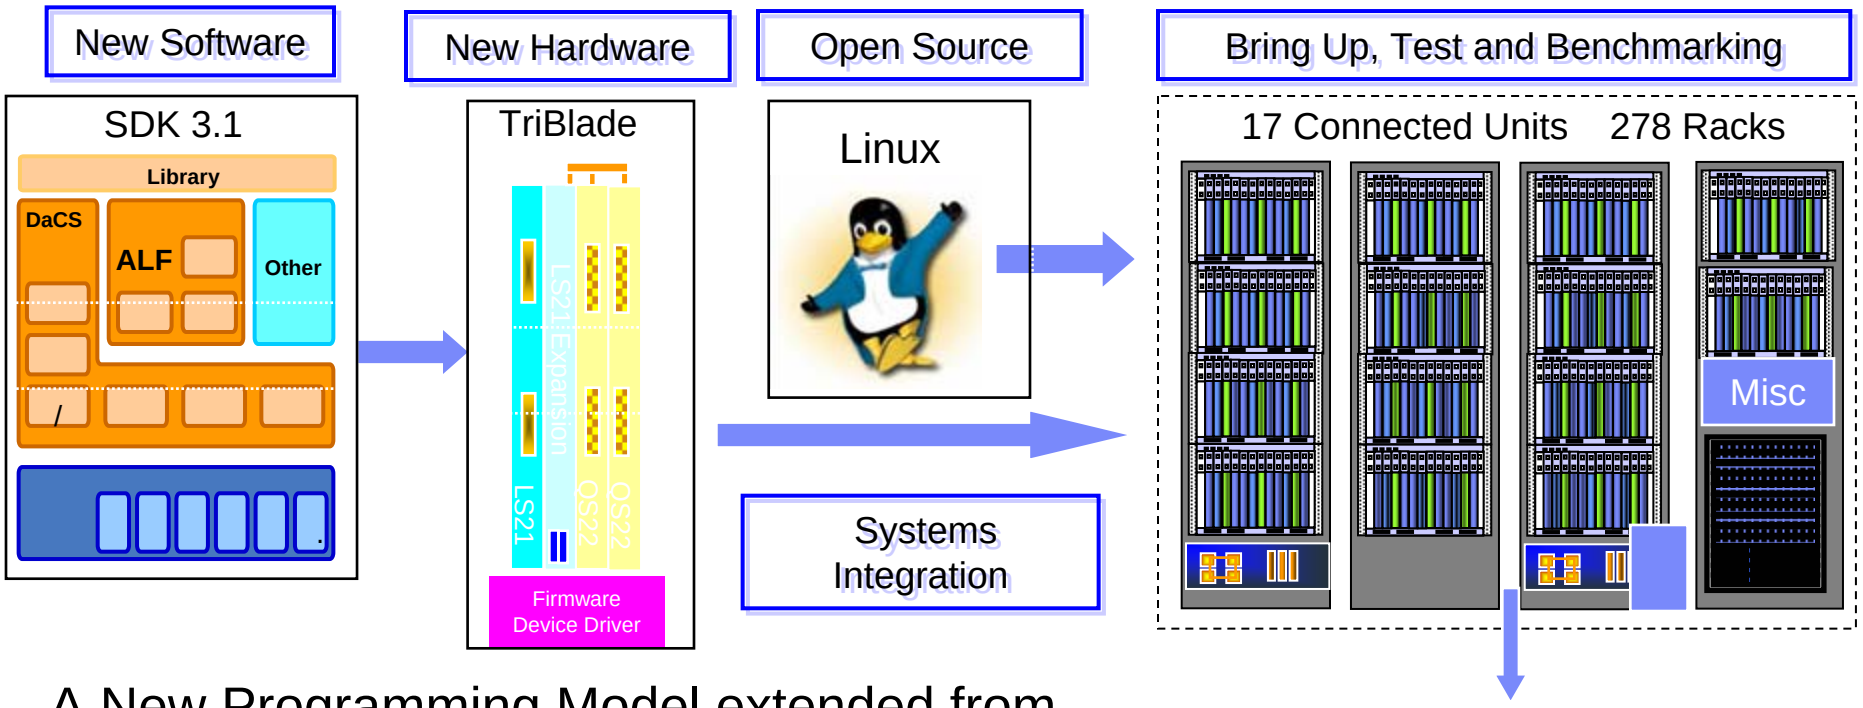

- A New Programming Model extended from standard, cluster computing (*Innovation derived through close collaboration with LANL*)
- Hybrid and Heterogeneous Hardware
- Built around BladeCenter and Industry IB-DDR

Role of Highly Scalable Multi-core Systems continues to Evolve

Blue Gene/P - a hallmark system for leadership performance in a space-saving, power-efficient package for the most demanding and scalable high-performance computing applications

System
72 Racks

Cabled 8x8x16

Rack
32 Node Cards
1024 chips, 4096 procs

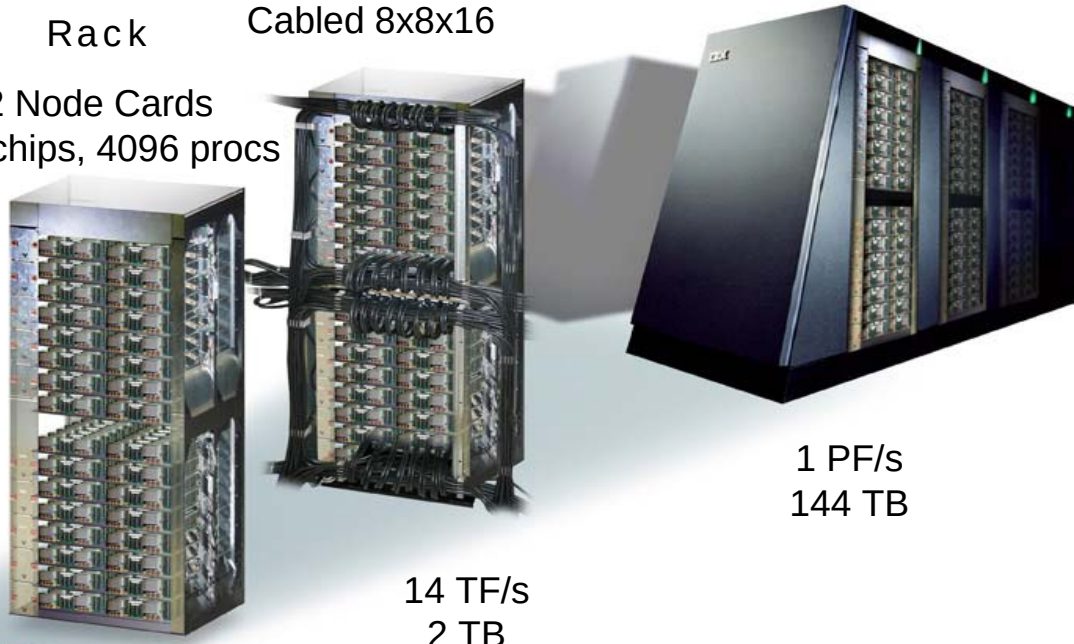

1 PF/s
144 TB

14 TF/s
2 TB

Node Card
(32 chips 4x4x2)
32 compute, 0-1 IO cards

435 GF/s
64 GB

Compute Card
1 chip, 20 DRAMs

13.6 GF/s
2.0 (or 4.0) GB DDR
Supports 4-way SMP

Chip
4 processors

13.6 GF/s
8 MB EDRAM

Front End Node / Service Node
POWER-based
blade & rack-based offerings

Highly Productive Computing Systems

Next-Generation POWER processor

High-speed Network

Scalable, Balanced System Design

Simple, Efficient Systems Mgmt.

Superb Reliability, Availability, Serviceability

Enhanced POWER Instruction Set Architecture

Advanced, Eco-Friendly Footprint, Packaging, Power & Cooling

Architecturally Robust Hardware & Software Design

DARPA HPCS Performance Goals:

Performance dimension	Improvement
Sustained execution speed	10X
Data processing rate	up to 50X
Memory access	up to 2,000X
Minimum bandwidth connecting system components	up to 3,000X

Five Decades of Eco-friendly System Innovations

S/360 Model 67 first virtualized machine

1960s -1970s

Air / liquid hybrid cooling technology

Thermal conduction cooling technology

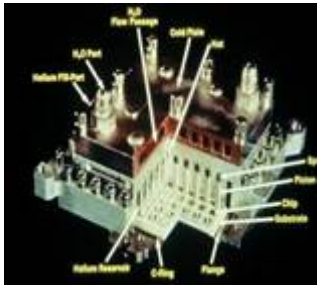

1980s

VM virtualization

CMOS processors

mid-1990s

Flat plate conduction cooling technology

Modular refrigeration cooling technology

late-1990s

Copper chip

High-k metal gates

IBM Energy Efficiency Institute, Austin, TX

2000s

Airgap

POWER6

Cell BE processor

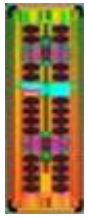

eDRAM

3D chip stacking

Holistic Design as Applied to Petascale Systems

Innovation from Atoms to Software

- **The simultaneous optimization of:**
  Materials, Devices, Circuits, Cores, Chips, System Architecture, System Assets and System Software
- **Provides the most effective means to optimize the value of IT offerings to the end user**
- **Catalyst for innovations in HPC:**
  Systems, Algorithms, Software Tools, Business Models

Note:

Execution relies upon the seamless integration of skills from across the spectrum

Market Dynamics and Design Principles; Petascale and Beyond

- Existing and emerging workloads with highly specialized (heterogeneous) requirements
- Increased societal awareness and regulatory mandate on environmentals/carbon footprint (power consumption, noise, space efficiency, resource lifecycle)
- Open, adaptive software ecosystem that supports legacy applications and hardware architectures while enabling new, advanced capabilities
- Robust, modular hardware architecture framework that enables new and existing processor types to be applied for workloads acceleration
- Reliability and Scalability for 10^{15} and beyond....

Next Era of Innovation – Hybrid Computing
The Next Bold Step in Innovation & Integration

Symmetric Multiprocessing Era

Today *pNext 1.0*

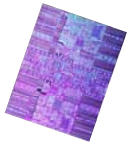
p6

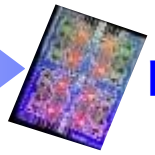
p7

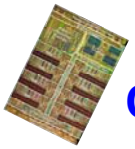
Cell


BlueGene

Technology Out

Driven by cores/threads

Hybrid Computing Era

pNext 2.0

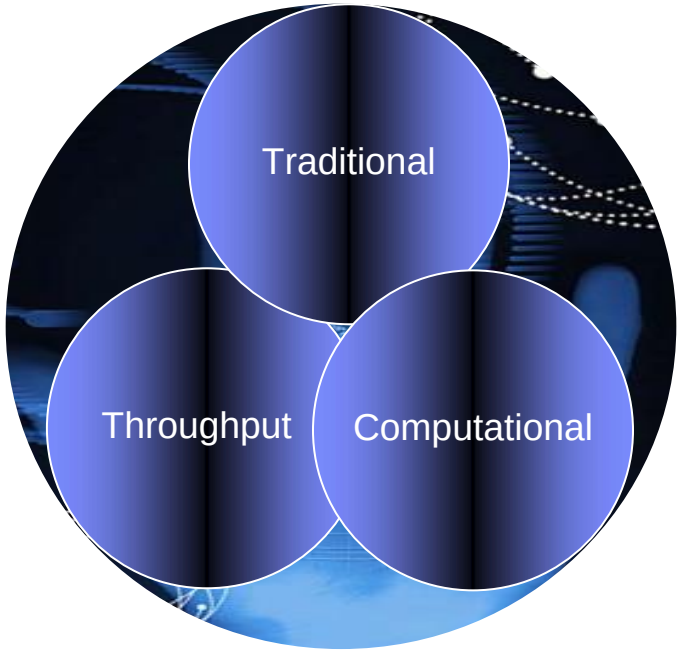

Market In

Driven by workload consolidation

HPC Cluster Directions

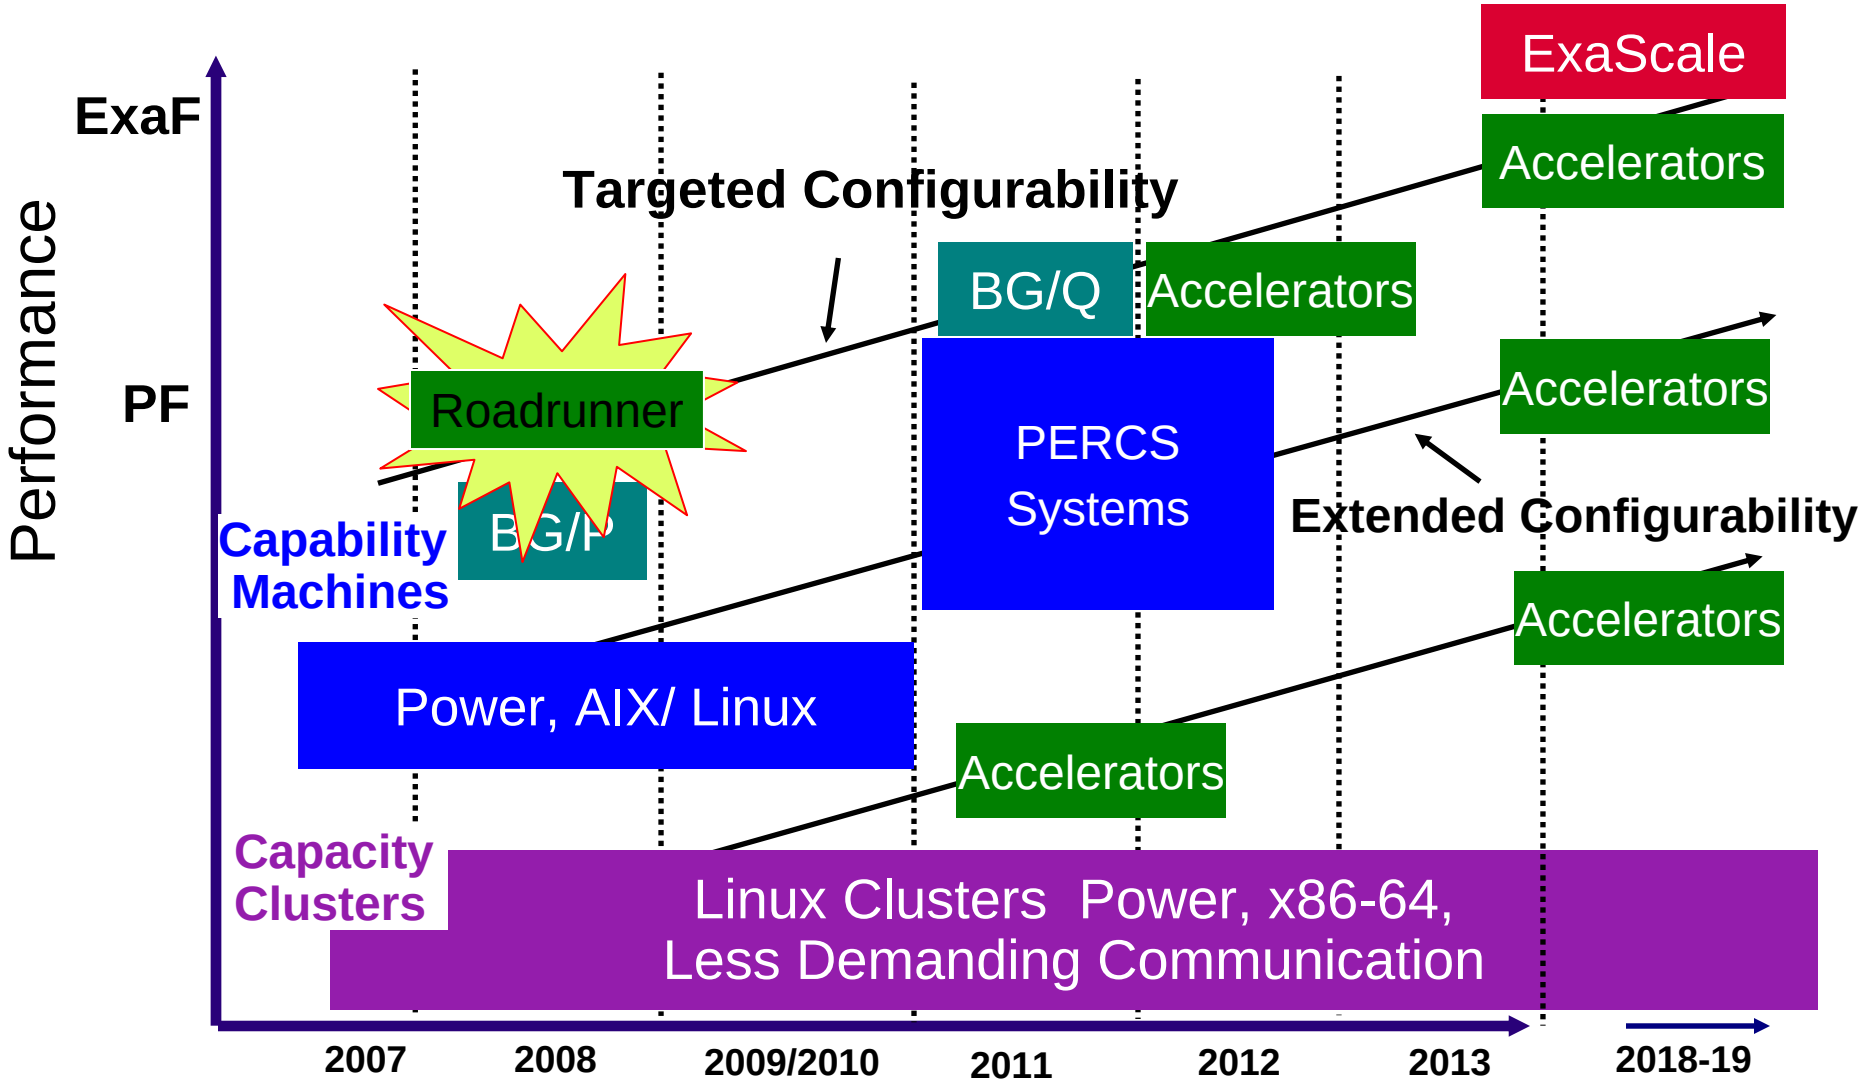

Programming Models: Architecture

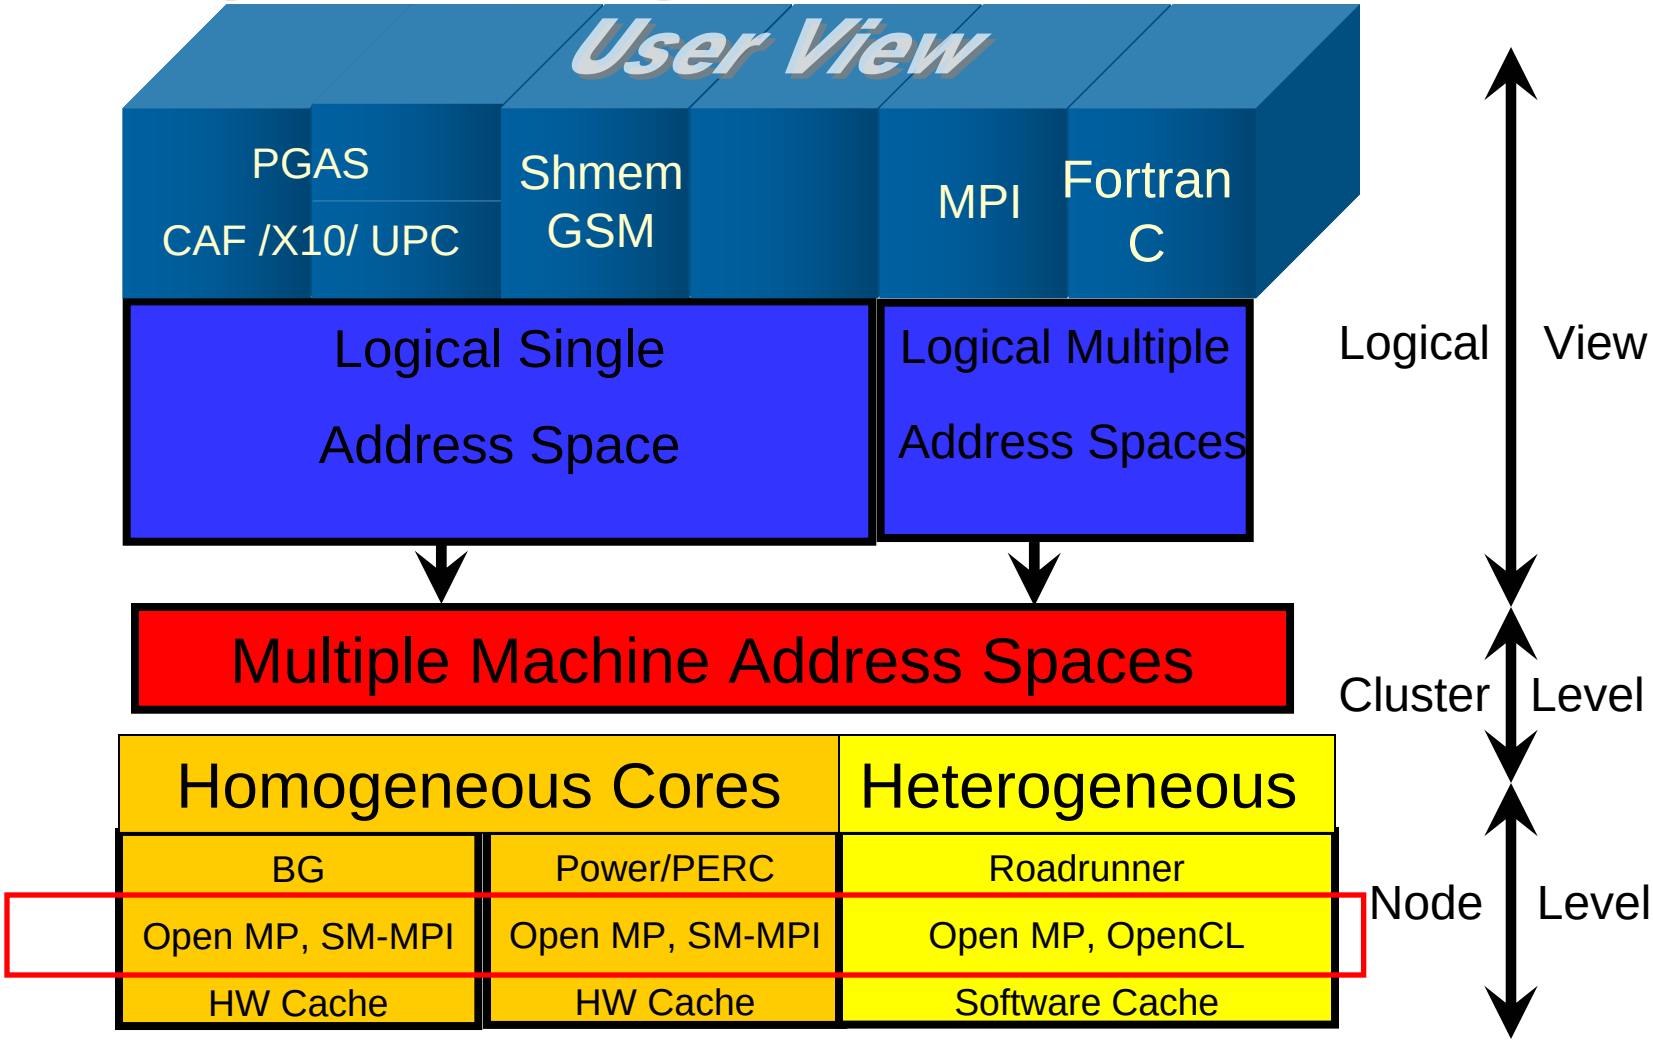

Two Standards

- Two standards evolving from different sides of the market

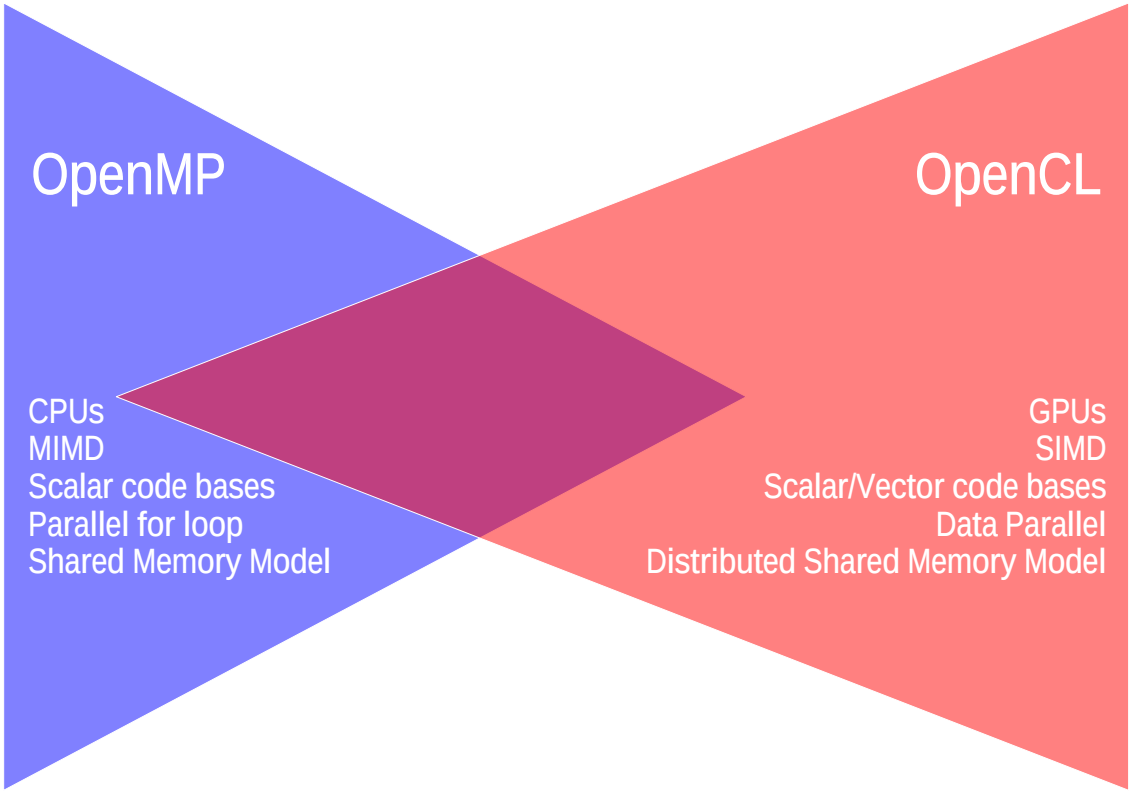

In the end, it's not about the technology; It's what you do with it that counts

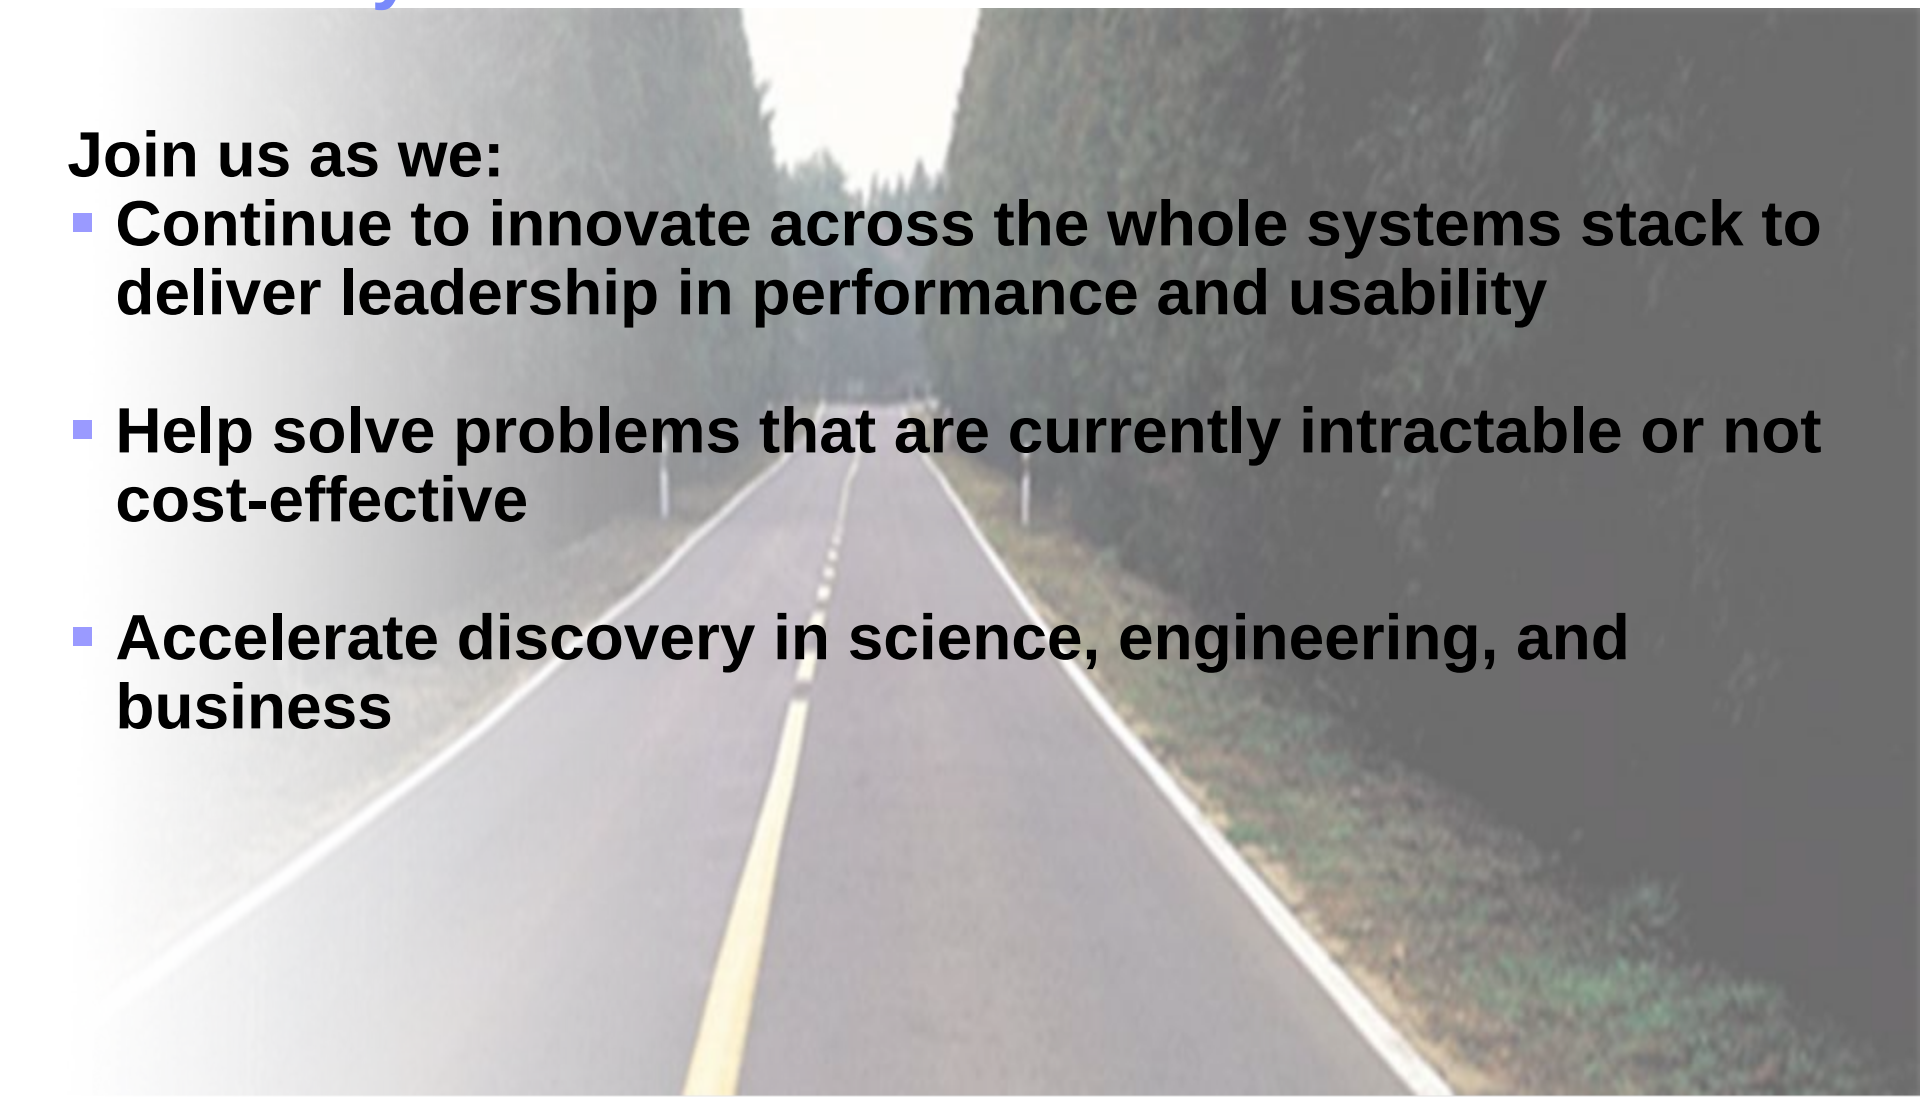

Join us as we:

- **Continue to innovate across the whole systems stack to deliver leadership in performance and usability**
- **Help solve problems that are currently intractable or not cost-effective**
- **Accelerate discovery in science, engineering, and business**

Thank you...

...any Questions?

Footnotes

Special notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Special notices (cont.)

The following terms are registered trademarks of International Business Machines Corporation in the United States and/or other countries: AIX, AIX/L, AIX/L(logo), alphaWorks, AS/400, BladeCenter, Blue Gene, Blue Lightning, C Set++, CICS, CICS/6000, ClusterProven, CT/2, DataHub, DataJoiner, DB2, DEEP BLUE, developerWorks, DirectTalk, Domino, DYNIX, DYNIX/ptx, e business(logo), e(logo)business, e(logo)server, Enterprise Storage Server, ESCON, FlashCopy, GDDM, i5/OS, IBM, IBM(logo), ibm.com, IBM Business Partner (logo), Informix, IntelliStation, IQ-Link, LANStreamer, LoadLeveler, Lotus, Lotus Notes, Lotusphere, Magstar, MediaStreamer, Micro Channel, MQSeries, Net.Data, Netfinity, NetView, Network Station, Notes, NUMA-Q, OpenPower, Operating System/2, Operating System/400, OS/2, OS/390, OS/400, Parallel Sysplex, PartnerLink, PartnerWorld, Passport Advantage, POWERparallel, Power PC 603, Power PC 604, PowerPC, PowerPC(logo), Predictive Failure Analysis, pSeries, PTX, ptx/ADMIN, RETAIN, RISC System/6000, RS/6000, RT Personal Computer, S/390, Scalable POWERparallel Systems, SecureWay, Sequent, ServerProven, SpaceBall, System/390, The Engines of e-business, THINK, Tivoli, Tivoli(logo), Tivoli Management Environment, Tivoli Ready(logo), TME, TotalStorage, TURBOWAYS, VisualAge, WebSphere, xSeries, z/OS, zSeries.

The following terms are trademarks of International Business Machines Corporation in the United States and/or other countries: Advanced Micro-Partitioning, AIX 5L, AIX PVMe, AS/400e, Chiphopper, Chipkill, Cloudscape, DB2 OLAP Server, DB2 Universal Database, DFDSM, DFSORT, DS4000, DS6000, DS8000, e-business(logo), e-business on demand, eServer, Express Middleware, Express Portfolio, Express Servers, Express Servers and Storage, General Purpose File System, GigaProcessor, GPFS, HACMP, HACMP/6000, IBM TotalStorage Proven, IBMLink, IMS, Intelligent Miner, iSeries, Micro-Partitioning, NUMACenter, On Demand Business logo, POWER, PowerExecutive, Power Architecture, Power Everywhere, Power Family, Power PC, PowerPC Architecture, PowerPC 603, PowerPC 603e, PowerPC 604, PowerPC 750, POWER2, POWER2 Architecture, POWER3, POWER4, POWER4+, POWER5, POWER5+, POWER6, POWER6+, pure XML, Redbooks, Sequent (logo), SequentLINK, Server Advantage, ServeRAID, Service Director, SmoothStart, SP, System i, System i5, System p, System p5, System Storage, System z, System z9, S/390 Parallel Enterprise Server, Tivoli Enterprise, TME 10, TotalStorage Proven, Ultramedia, VideoCharger, Virtualization Engine, Visualization Data Explorer, X-Architecture, z/Architecture, z/9.

A full list of U.S. trademarks owned by IBM may be found at: http://www.ibm.com/legal/copytrade.shtml.
The Power Architecture and Power.org wordmarks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org.
UNIX is a registered trademark of The Open Group in the United States, other countries or both.
Linux is a trademark of Linus Torvalds in the United States, other countries or both.
Microsoft, Windows, Windows NT and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries or both.
Intel, Itanium, Pentium are registered tradas and Xeon is a trademark of Intel Corporation or its subsidiaries in the United States, other countries or both.
AMD Opteron is a trademark of Advanced Micro Devices, Inc.
Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries or both.
TPC-C and TPC-H are trademarks of the Transaction Performance Processing Council (TPPC).
SPECint, SPECfp, SPECjbb, SPECweb, SPECjAppServer, SPEC OMP, SPECviewperf, SPECapc, SPEChpc, SPECjvm, SPECmail, SPECimap and SPECsfs are trademarks of the Standard Performance Evaluation Corp (SPEC).
NetBench is a registered trademark of Ziff Davis Media in the United States, other countries or both.
AltiVec is a trademark of Freescale Semiconductor, Inc.
Cell Broadband Engine is a trademark of Sony Computer Entertainment Inc.
Other company, product and service names may be trademarks or service marks of others.