# IBM Systems and Technology Group University 2005

San Diego, California
January 16 - 19, 2005

ON DEMAND BUSINESS™

v 1.0

© 2005 IBM Corporation

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

IBM Systems and Technology Group University 2005

# Deep Computing Innovation: New Offerings For 2005

Course #: CB52

Rebecca Austen
Director, Deep Computing Marketing


v 1.0

© 2005 IBM Corporation

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Agenda

- **Deep Computing Overview**
  - Application, Technology and Market Trends
  - Deep Computing Strategy and Portfolio
- **Deep Computing Innovation**
  - Blue Gene
  - Cell
  - Deep Computing Visualization
  - Deep Computing Capacity on Demand
- **Conclusion**
- **Resources**

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Learning Objectives

- Understand the Deep Computing market place
- Identify Deep Computing opportunities in your industry
- Evaluate emerging Deep Computing technologies and solutions which enable your clients to gain new insights and solve their most challenging problems
- Know whom to contact and where to find more information on Deep Computing

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Deep Computing Overview

- **Application, Technology and Market Trends**
- **Deep Computing Strategy and Portfolio**

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# What Drives HPC? --- *"The Need for Speed…"*

**Computational Needs of Technical, Scientific, Digital Media and Business Applications Approach or Exceed the Petaflops/s Range**

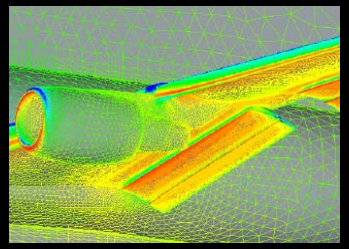

### CFD Wing Simulation

512x64x256 Grid
( 8.3 x 10e6 mesh points)
5000 FLOPs per mesh point,
5000 time steps/cycles

**2.15 x 10e14 FLOPs**

Source: A. Jameson, et al

### CFD Full Plane Simulation

512x64x256 Grid
( 3.5 x 10e17 mesh points)
5000 FLOPs per mesh point
5000 time steps/cycles

**8.7x 10e24 FLOPs**

### Materials Science

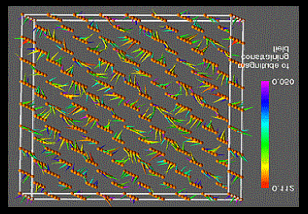

Magnetic Materials:

Current: 2000 atoms; 2.64 TF/s, 512GB

Future: HDD Simulation – **30TF/s, 2 TBs**

Electronic Structures:

Current: 300 atoms; **0.5 TF/s, 100GB**

Future: 3000 atoms; **50TF/s, 2TB**

Source: D. Bailey, NERSC

### Digital Movies and Special Effects

~ 1E14 FLOPs per frame
50 frames/sec
90 minute movie

- 2.7E19 **FLOPs**
-

~ 150 days on 2000
1 GFLOP/s CPUs

Source: Pixar

### Spare Parts Inventory Planning

Modeling the optimized deployment of 10,000 part numbers across 100 parts depots and requires:

- 2 x 10e14 FLOP/s
  ( 12 hours on 10, 650MHz CPUs)
- 2.4 PetaFlop/s sust. performance
  ( 1 hour turn-around time )

Industry trend for rapid, frequent modeling for timely business decision support drives higher sustained performance

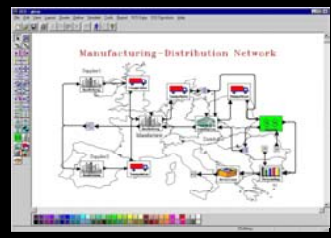


Source: B. Dietrich, IBM

# Deep Computing: Energy Production and Trading

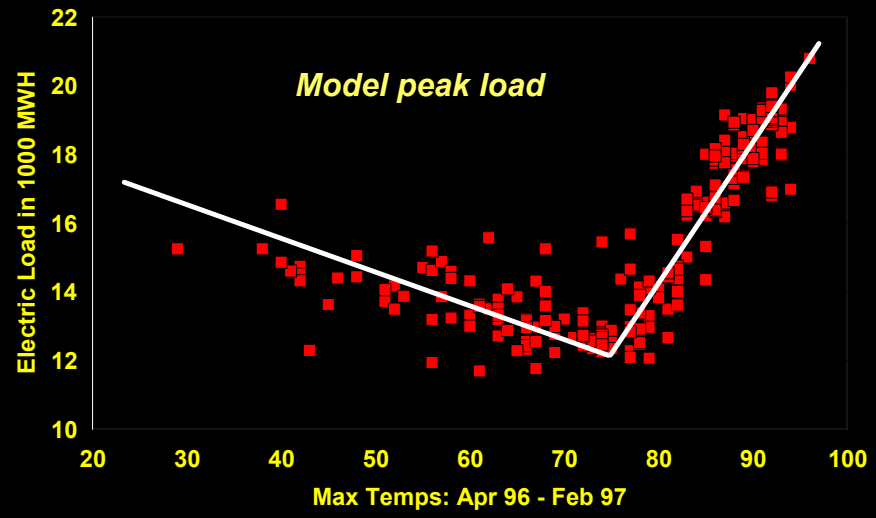


The application of computational methods
- •Mathematical modeling
- •Optimization
- •Statistics
- •Design and analysis of algorithms
- •Parallel computing

to large data sets
- •Historical
- •Operational

to solve business decision problems

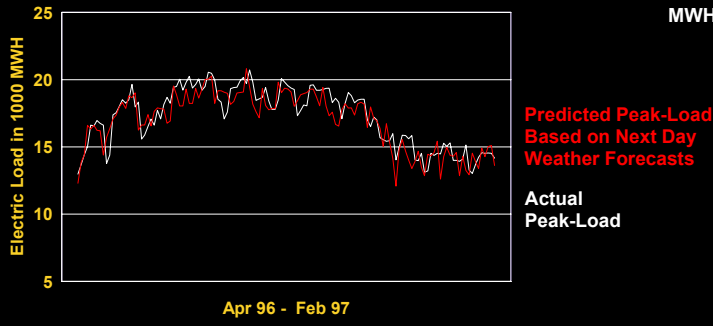


*Predict peak load based on weather forecast*

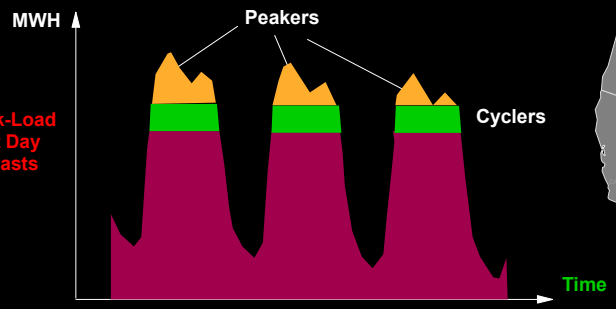


*Plan production schedules*

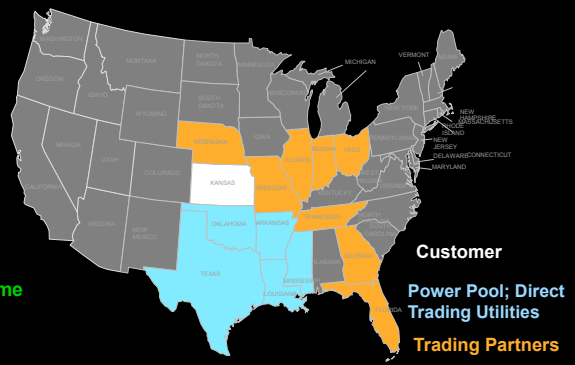


*Identify energy trading partners*

Source: W. Pulleyblank, IBM

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# LOFAR (Low Frequency Array)

**An array of radio antennae covering 300 km in Northern Europe and sending 320 Gbits/s of data back to a central computer for processing**

LOFAR will observe galaxies at the edge of the visible universe, revealing how they formed, 13 billion years ago.

Data streaming: also useful for surveillance, weather, and geological sensing.

Will ultimately require multi-petaflop processing

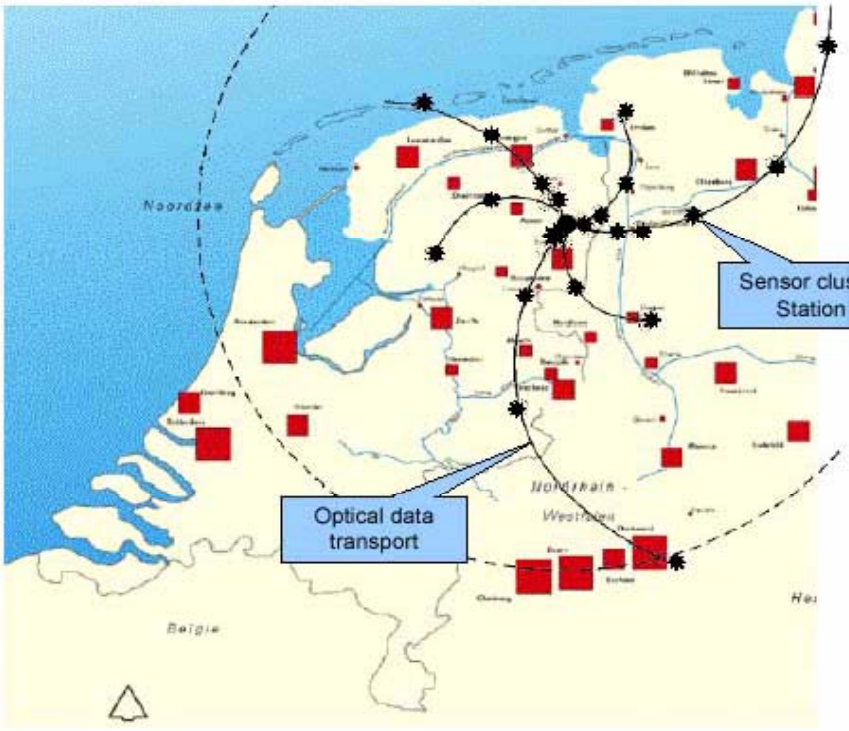


Proposed array in Netherlands

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Market Opportunity & Growth

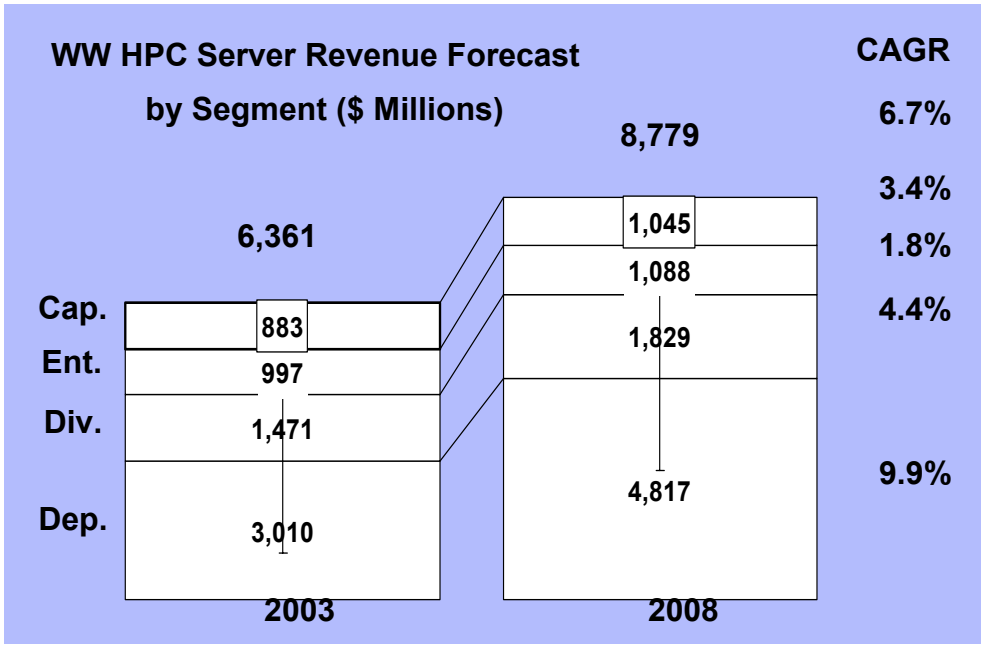


Source: IDC Cluster Multi-Client Study, May 2004

- Total HPC server market growth is attractive at 6.7%
- Departmental segment (<$250K) has almost 3X the revenue and 2X the growth rate of next largest segment (Divisional)
- Departmental & Divisional (<$1M) represent 70% of 2003 revenue (growing to 75% in 2008)
- Total HPC market opportunity (servers, storage, workstations, software, etc.) estimated at $32B to $40B
- Strong growth in Linux clusters and in emerging business areas such as life sciences, digital media, and financial analytics

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# HPC Server Revenue and Share

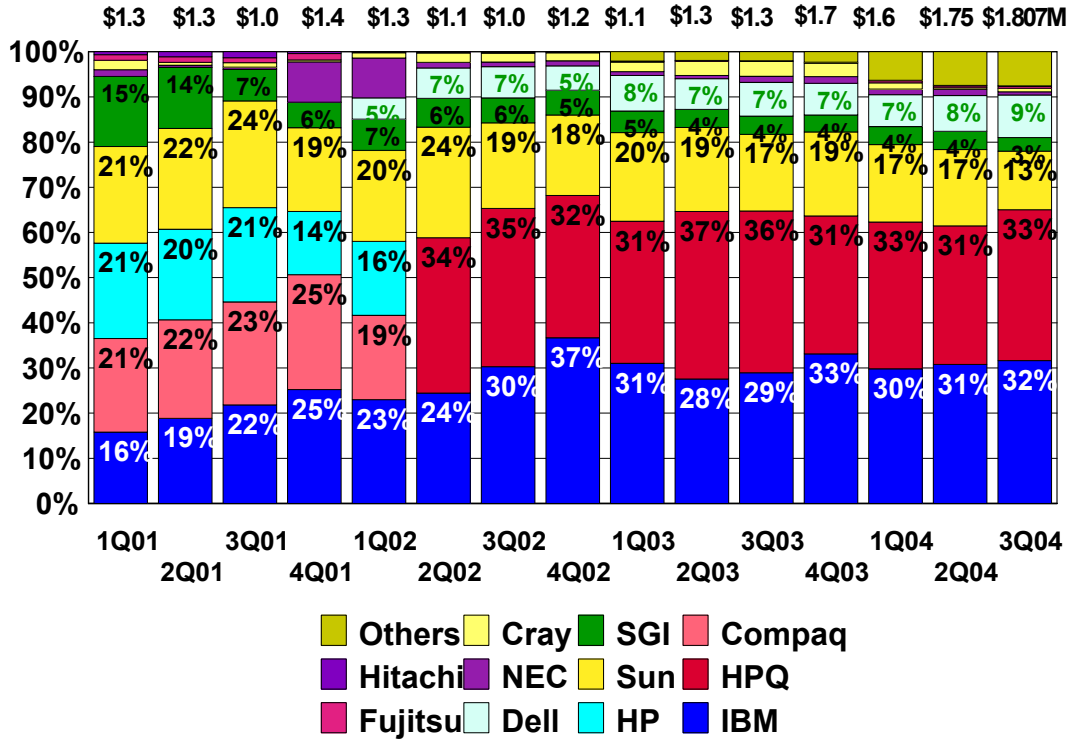


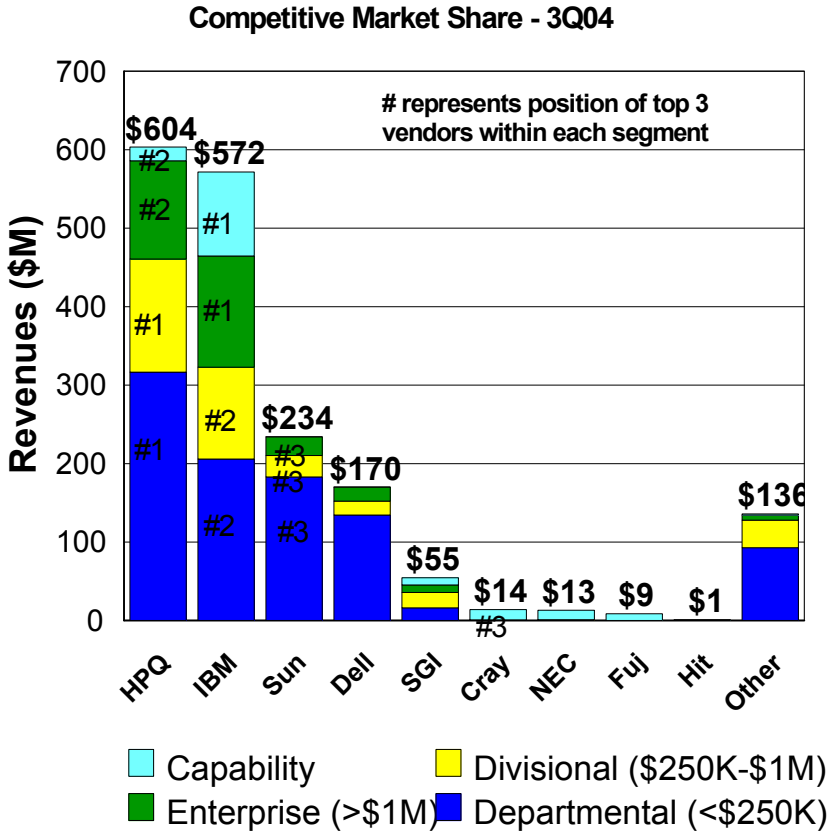


Source: IDC 12/2004

# Top500 List of Supercomputers

| | *Ten years ago* | *Five years ago* | *Today* |
|---|---|---|---|
| *Largest system* | 143 Gflops | 2.1 Tflops | 70.7 Tflops |
| *Teraflop systems* | 0 | 2 | 398 |
| *Research/Academic* | 60% | 48% | 41% |
| *Industry* | 24% | 46% | 55% |
| *Linux clusters* | 0 | 6 | 294 |

## IBM Leadership (Nov 2004)

- #1 System – DOE - BlueGene/L (70.7 TF)
- Most entries on list with 216 (43.2%)
- Most aggregate throughput 556.9 TF (49.4%)
- Most in Top10 (4), Top20 (8), Top100 (58)
- Largest system in Europe (MareNostrum)
- Most Linux Clusters with 161 of 294 (54.7%)

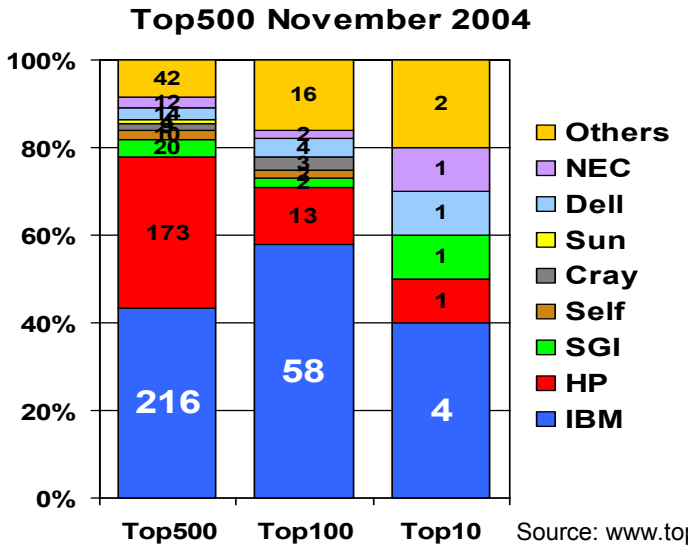


Source: www.top500.org

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# The Evolution of Large Scale Computing

64w SMP

SCALE UP

**The Evolution to "Scale-Out" from "Scale-Up" Computing was Driven to Lower the Cost of Computing**

1U/1-2w

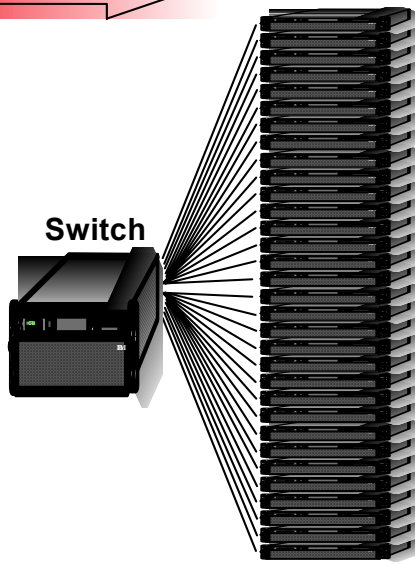

Switch



SCALE OUT

2U/8-16w

Switch

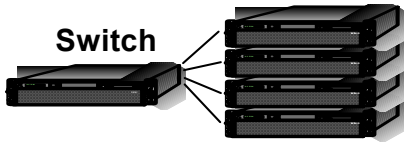

SCALE IN

**"Scale-In" Computing Is the Next Logical Step:**

- **Equal Server Price**
- **Reduced Switch Cost**
- **Higher Inter-CPU BW**
- **Improved Simplicity**
- **Higher Compute Density**
- **Better Environementals**
- **Better System Reliability due to Fewer Components**

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Hybrid delivery models are emerging as result of shifting economics

- **Variances in workload requirements (“peaks and valleys”)**
- **Software licensing costs not keeping up with Moore’s Law**
- **Need to focus on core business/research competencies (not IT)**

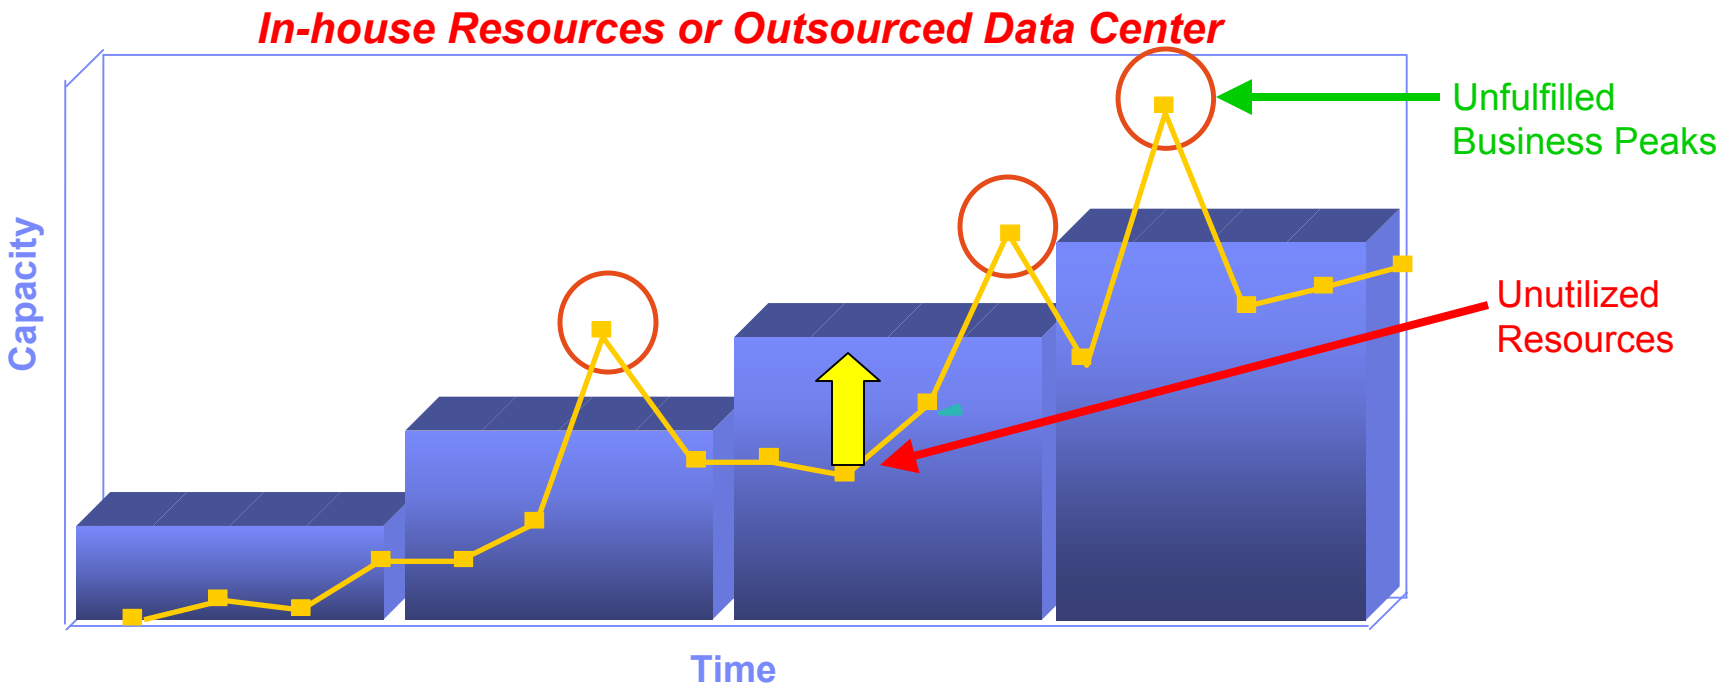

# New challenges and disruptions are stimulating innovation

- **Environmental issues are driving new system design concepts**
  - Power, cooling and space constraints → extreme density and scalability at lower frequencies
- **Data management surfaces as critical infrastructure requirement**
  - Archival and retention → sophisticated data models
  - Regulatory compliance → industry-tailored content management
  - Simulation of data → advanced visualization systems
- **Global security and defense requirements growing in complexity**
  - Capture and analysis of massive quantities of digital data → sophisticated business intelligence and data mining techniques, digital content management

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Deep Computing

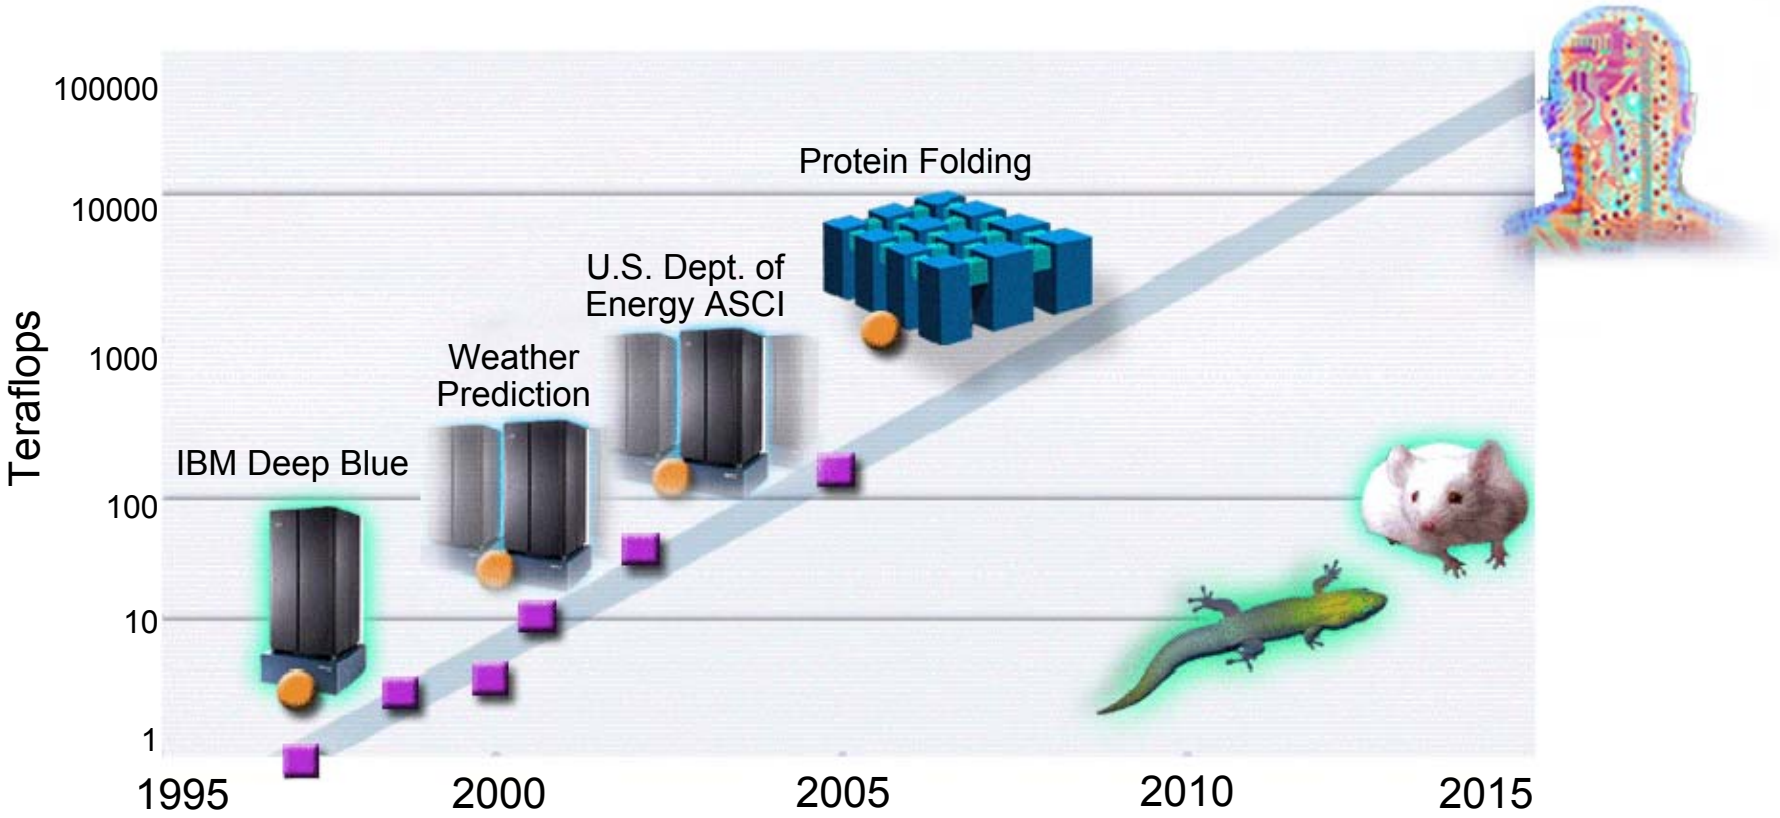


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Deep Computing Leadership Initiatives

- **Helping solve mission critical, computationally-intensive problems**
- **Extending the amount of science and engineering that can be supported by available computational resources**
- **Enabling global business competitiveness by devising solutions to industry's most complex and challenging problems, generating strategic value**
- **Developing deep expertise in all parts of the value chain**

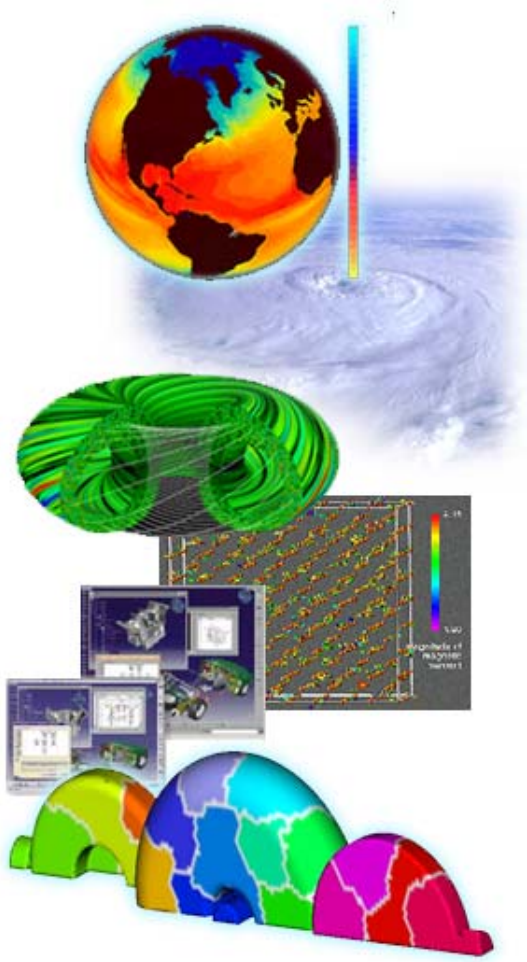

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# IBM's Deep Computing Strategy

## *Solving Problems More Quickly at Lower Cost*

- **Aggressively evolve and improve POWER-based Deep Computing product line**
- **Develop advanced systems based on loosely coupled clusters**
- **Research and overcome obstacles to parallelism and other revolutionary approaches to supercomputing**
- **Increase means of accessing supercomputing with on demand capabilities**

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Deep Computing Embraces a Broad Spectrum of Markets

## Digital Media

**Digital content creation, management and distribution**

## Life Sciences

**Research, drug discovery, diagnostics, information-based medicine**

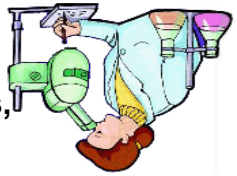

## Business Intelligence

**Data warehousing and data mining**

## Petroleum

**Oil and gas exploration and production**

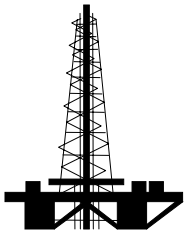

## Financial Services

**Optimizing IT infrastructure, risk management and compliance, analytics**

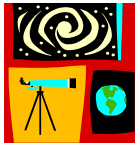

## Industrial/Product Lifecycle Management

**CAE, EDA, CAD/PDM for electronics, automotive, and aerospace**

## Government & Higher Education

**Scientific research, classified/defense, weather/environmental sciences**

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# IBM Systems – Industry Leadership and Choice

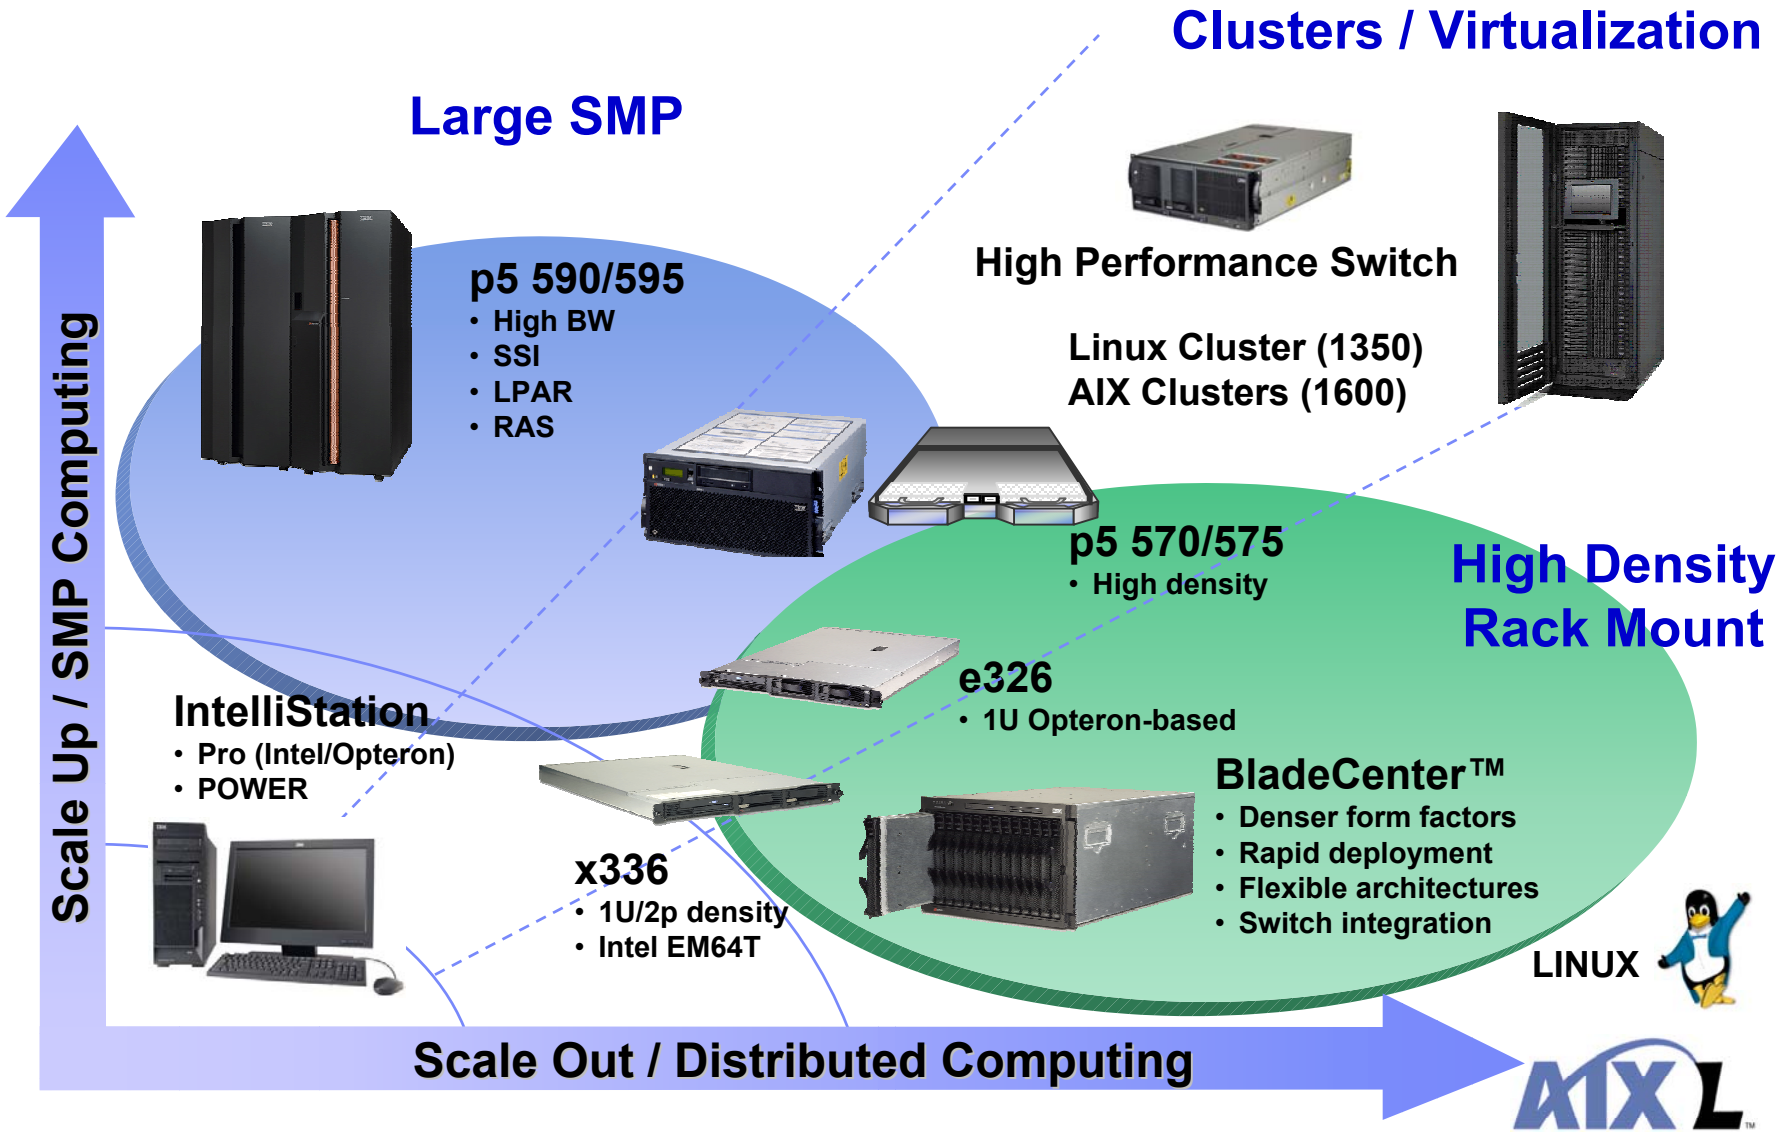


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Deep Computing Portfolio – A Breadth of Choices

- ***System Hardware***
  - Servers, blades, clusters & workstations
    - POWER, Intel, Opteron
  - Visualization/rendering
  - High performance interconnects
    - e.g., HPS, Myrinet, Infiniband
  - Storage systems, storage virtualization
- ***System Software***
  - Linux, Unix, Windows
  - System & storage management
    - e.g., CSM, PSSP, GPFS
  - DB2, WebSphere, Tivoli
  - Grid & on demand middleware
- ***Special-purpose systems***
  - Blue Gene/L
  - Gov't & research partnerships (e.g., ASCI)
- ***Applications & Tools***
  - ISV & in-house software
  - Compilers, schedulers, libraries, tools
    - 3rd party, ESSL/PESSL, LoadLeveler
  - Open source & public domain codes
  - ACTC tools
- ***Services***
  - IGS services practices & consulting
  - Hosting & utility services
  - IBM Global Financing
- ***Solutions***
  - Deep Computing Capacity on Demand
  - Visualization solutions
  - Infrastructure solutions (e.g., Grid)
  - Enterprise application solutions
  - Industry-specific solutions

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Deep Computing Innovation

- **Blue Gene**
- **Cell**
- **Deep Computing Visualization**
- **Deep Computing Capacity on Demand**

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# IBM eServer Blue Gene Solution

**Leadership HPC performance**

**Broad applicability to important HPC workloads helps advance science, engineering and business**

**Enhanced accessibility to world-class computing**

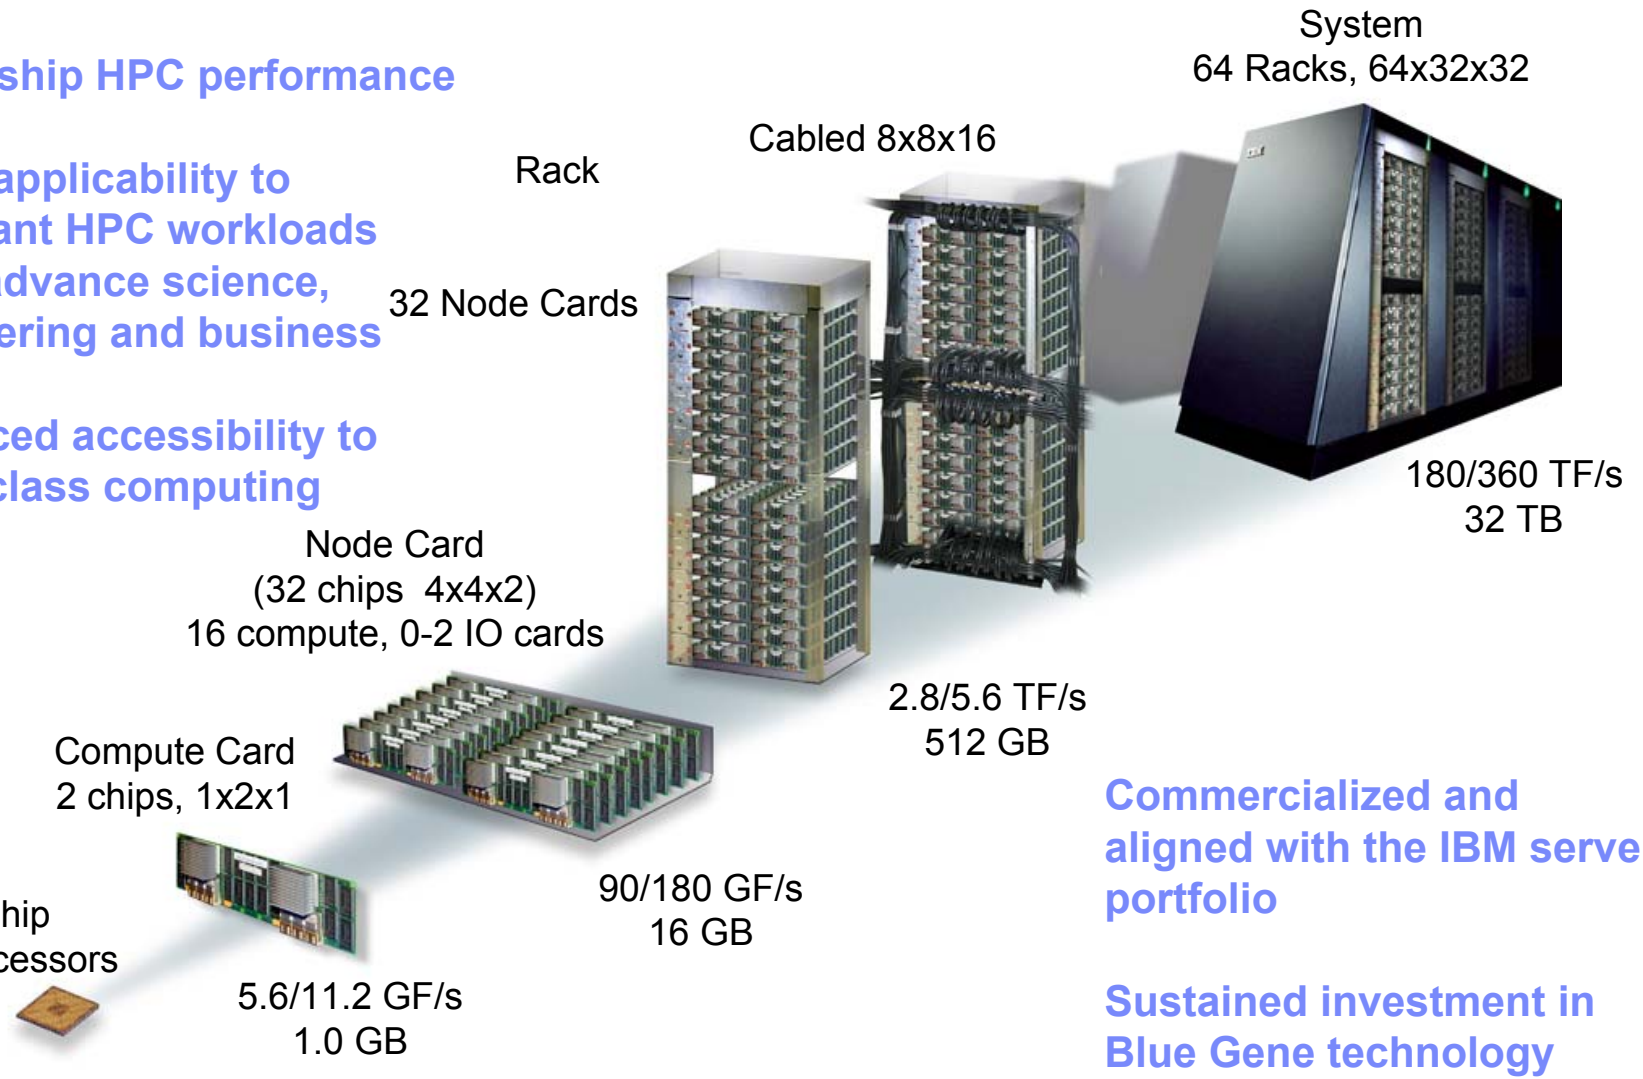


**Commercialized and aligned with the IBM server portfolio**

**Sustained investment in Blue Gene technology**

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# A Brief History of the Blue Gene Program

- **December 1999, IBM Research announces a 5 year, $100M effort to build a petaflop scale supercomputer to apply to complex problems such as protein folding**
- **The Blue Gene project starts with two primary goals:**
  - Advance the state of the art of biomolecular simulation.
  - Advance the state of the art in computer design and software for extremely large scale systems.
- **November 2001, a research partnership with Lawrence Livermore National Laboratory is announced.**
  - The suite of potential applications was augmented to a broad range of deep computing problems.
- **November 2002, the acquisition of a BG/L machine by LLNL from IBM is announced as part of the ASCI Purple contract.**
- **June 2003, first chips arrived and powered on**
- **June 2004, two BG/L systems are among the world's 10 fastest computers as rated by the TOP500 organization**
- **November 2004, Blue Gene is the #1 system on the TOP500 – Commercialization plans announced**


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# BlueGene Value Proposition

***For:*** **Scientists, Researchers, Engineers and Principal Investigators in selected HPC application segments**

***Who Need:*** **Extremely high, scalable computational performance (100's of TFLOPS) balanced with efficient environmental characteristics**

***IBM Provides:*** **The eServer Blue Gene Solution, a dense, highly scalable supercomputer optimized for HPC parallel programming workloads**

***Unlike:*** **Cray, NEC and SGI**

***IBM's Offering Has:*** **The most densely packaged computational engine, delivering unsurpassed scalability and offering the best ratios of performance to power, cooling and floor space consumed**

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Market Segmentation and Selection Based on Overall HPC Market Opportunity

| | Application Segment | CAGR 2003-2006 | Proportion of HPC Opportunity | Affinity to BlueGene Architecture | Relative Profitability | Overall Assessment of Attractiveness |
|---|---|---|---|---|---|---|
| | Scientific Research | Low (3.9%) | High (22.6%) | High | Low | High |
| | Biosciences | High (10.3%) | High (18.7%) | High | Medium | High |
| | Classified Defense | Medium (4.7%) | High (13.3%) | High | Low | High |
| Non Public Sector Opportunities | CAE | Medium (6.8%) | High (16.3%) | Medium | High | High |
| | Chemical Eng'g | High (9.1%) | Low (3.6%) | Medium | Medium | Medium |
| | Petroleum | Medium (5.8%) | Medium (6.3%) | Medium | Medium | Medium |
| | FSS | High (8.7%) | Low (3.5%) | Medium | High | Medium |
| | EDA | Medium (7.1%) | Medium (6.9%) | Medium | Medium | Medium |
| | Digital Media | Medium (7.2%) | Low (1.5%) | Medium | High | Low |
| | PDM | Low (3.8%) | Low (4.0%) | Low | Medium | Low |


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Blue Gene – Likely Best-Matched Applications

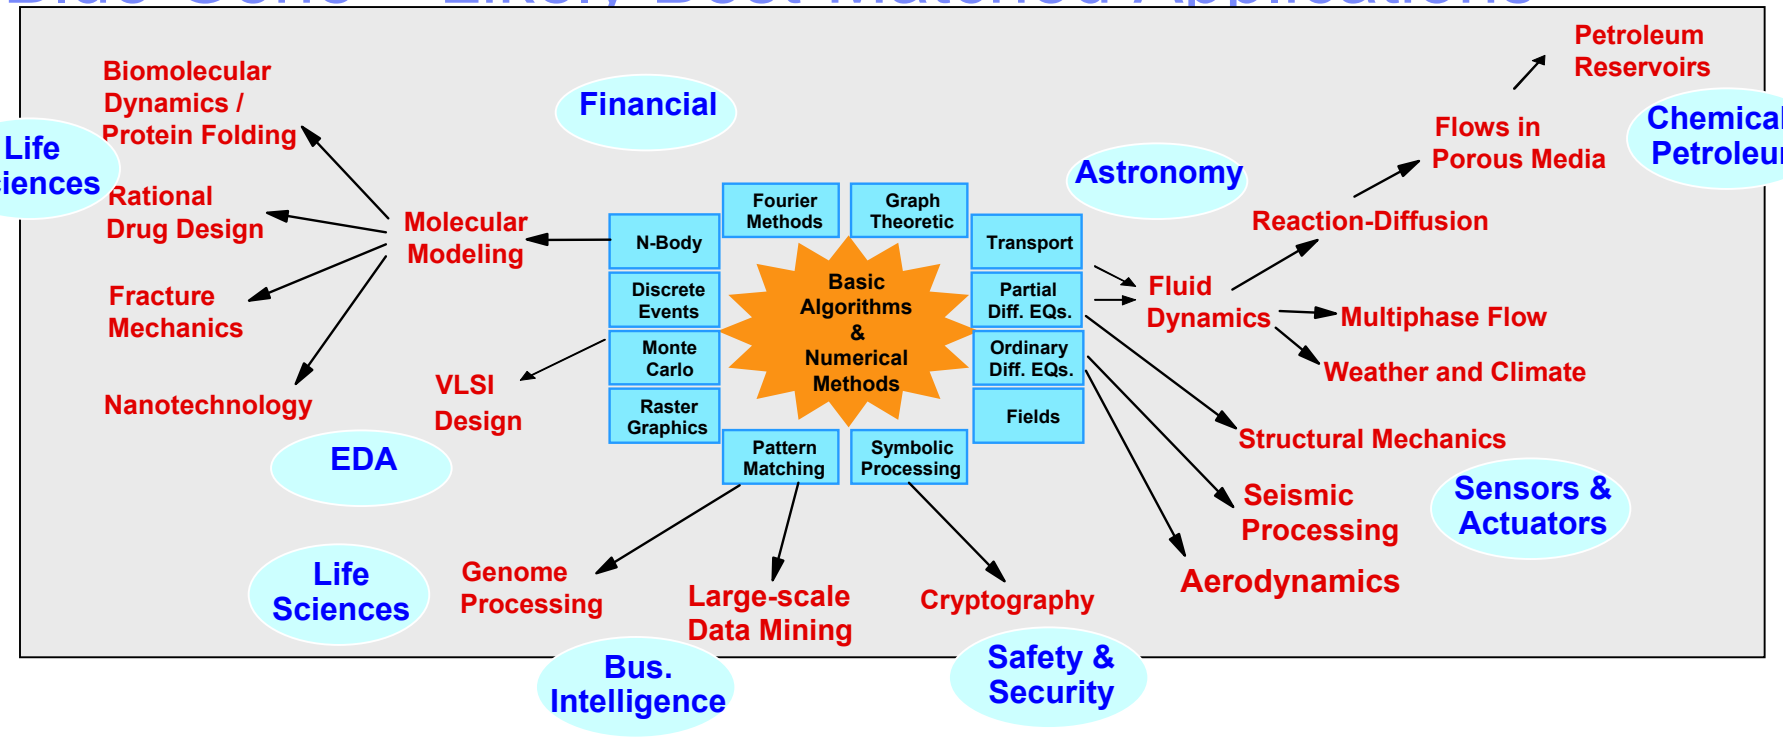


- Best-matched apps: **Intense interactions between 1,000s to 1,000,000s of simple units, with high ratios of Flops & cluster BW vs. memory size & disk BW**
  - Physical simulation (chemistry, biology, aerodynamics, etc.), e.g., atoms (Protein Folding/Drug Design), air pockets (Weather and Climate), logic gates (VLSI simulation), etc.
  - High inherent parallelism, tight coupling between many elements, with neighbor communication.
- Also matched: **Purely-parallel searches over huge data or parameter spaces in memory**
  - Genome searching, Cryptography, Data mining, etc.
- Non-Candidate apps: **Client/server structures; Dominated by local disk I/O; Larger memory footprint local processes**


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# BlueGene System Configuration and Business Model

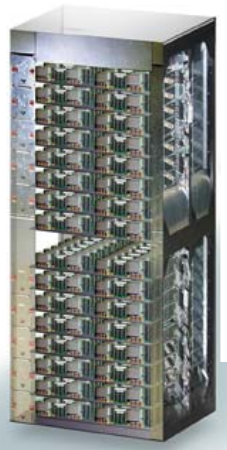

**Node Card**
Up to 32 chips (4x4x2)
Up to 16 Compute Nodes
Up to 2 IO Nodes
Up to 90/180 GF/s
Up to 16 GB Memory

**Compute Card**
2 chips (1x2x1)
5.6/11.2 GF/s
1 GB Memory

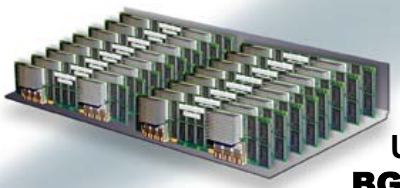

**BG/L Chip**
2 processors
2.8/5.6 GF/s
4 MB L2

**BG/L Core Rack**
Up to 32 Node Cards
Up to 2.8/5.6 TF/s
Up to 512 GB Memory

**BG/L Program Software**
CNK, SMCS, MCP, GNU Tool

**Host System**
Service Node (P4/SLES8)
Front End Nodes (P5/SLES9)
Storage System, Ethernet Switch, Cabling, DB2, SuSE, Java, Compilers, LL, GPFS

IBM Deliverables under E&TS SOW

Customer Deliverables under PO to SG / S&D

**BlueGene/L Deal**



This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# BlueGene/L Software Hierarchical Organization

- **Compute nodes dedicated to running user application, and almost nothing else - simple compute node kernel (CNK)**
- **I/O nodes run Linux and provide a more complete range of OS services – files, sockets, process launch, signaling, debugging, and termination**
- **Service node performs system management services (e.g., heart beating, monitoring errors) - transparent to application software**

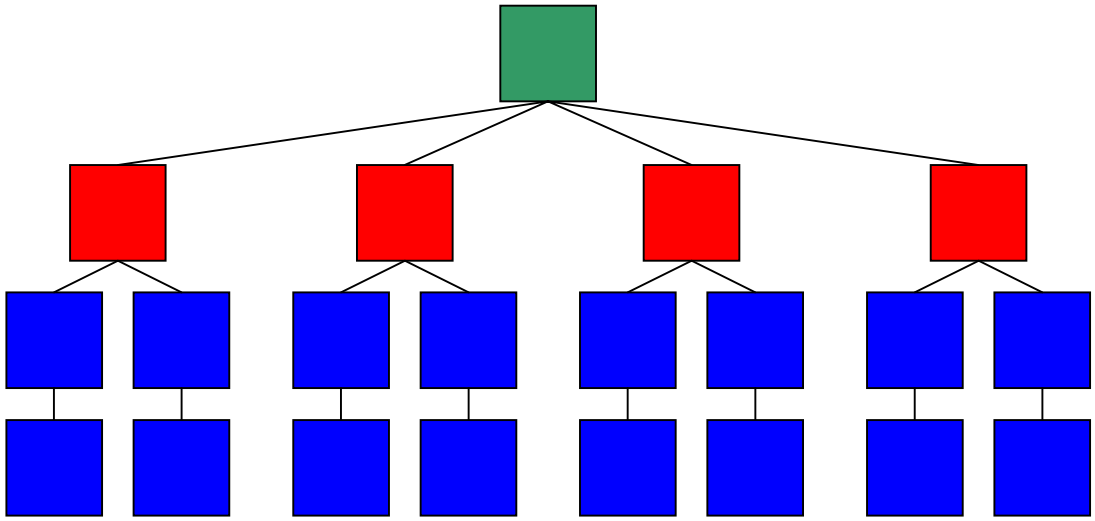


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Programming Environment and Scalability

- **Familiarity**
  - Standard programming languages, parallel computing model - Fortran, C, C++ with MPI
    - Full language support, automatic SIMD FPU exploitation by IBM compilers
  - Linux based development environment
    - User interacts with system through FE nodes running Linux – compilation, job submission, debugging
    - Compute Node Kernel provides look and feel of a Linux environment – POSIX system calls (with restrictions to ensure scaling)
- **Scalability**
  - Hierarchical organization of familiar functionality
    - Linux OS on I/O nodes, system management services on Service node
  - MPI implementation – supports standards, but designed for scalability to ~100K processors
  - Simplicity – used to drive both performance and reliability (e.g., partitioning only for space sharing, no time sharing; constraint on memory size used per node)
- **Scaling Experiences**
  - A number of scientific applications have been scaled successfully to 8K processors, most often within 1 to 4 run attempts:
    - Ab initio molecular dynamics (Materials and Life Sciences)
    - Turbulence and instability modeling (Classified/Defense)
    - Protein dynamics (Life Sciences)
    - Lattice quantum chromodynamics (Physics Research)
    - Nuclear astrophysical simulations (Astronomy)


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# HPC Tools Software for Blue Gene (2005 Outlook)

## IBM Software Stack

- **XL Compilers**
  - Externals preserved
  - New options to optimize for specific Blue Gene functions
- **LoadLeveler**
  - Same externals for job submission and system query functions
  - Backfill scheduling to achieve maximum system utilization
- **GPFS**
  - Provides high performance file access, as in current pSeries and xSeries clusters
  - Runs on IO nodes and disk servers
- **ESSL/MASSV**
  - Optimization library and intrinsics for better application performance
  - Serial Static Library supporting 32-bit applications
  - Callable from FORTRAN, C, and C++
- MPI
  - Message passing library, based on MPICH2, tuned for Blue Gene architecture

## Other Software Support

- **Etnus TotalView**
  - –Parallel Debugger
- **Lustre File System**
  - –Enablement underway at LLNL
- **GNU Compilers**
  - –Ported to Blue Gene
- **FFT Library**
  - –Tuned functions by TU-Vienna
- **Performance Tools**
  - –HPC Toolkit
  - –Paraver


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Blue Gene Benefits – At a Glance

| Attribute | Details | Benefits |
|---|---|---|
| Processor | PowerPC 440 700MHz; two per node | Low power allows dense packaging; better processor-memory balance |
| Memory | 512 MB SDRAM-DDR per node | |
| Networks | 3D Torus - 175MB/sec in each direction<br>Collective Network – 350MB/sec; 1.5 usec latency<br>Global Barrier/Interrupt<br>Gigabit Ethernet (machine control and outside connectivity) | Special networks speed up internode communications; designed for MPI programming constructs; improve systems management |
| Compute Nodes | Dual processor; 1024 per rack | Double FPU improves performance |
| I/O Nodes | Dual processor; 16 per rack (additional nodes optional) | Strengthens systems management |
| Operating Systems | Compute Node – Lightweight proprietary kernel<br>I/O Node – Embedded Linux<br>Front End and Service Nodes – SuSE SLES 9 Linux | Kernel tailored to processor design; industry-standard distribution preserves familiarity to end user |
| Performance | Peak per rack (virtual node mode) – 5.73 teraflops<br>Peak per rack (coprocessor mode) – 2.86 teraflops<br>Linpack per rack (VN mode) – 4.53 teraflops | Highest available performance benefits capability customers |
| Power | 28.14 kW power consumption per rack (maximum)<br>208 VAC 3-phase; 100 amp service per rack | Low power draw enables dense packaging |
| Cooling | Air conditioning 8 tons/rack (minimum)<br>2800 CFM (compute rack); 350 CFM (power supplies) | Low cooling requirements enable extreme scale-up |
| Acoustics | 9.0 LwAD and 8.7 LwAm | |
| Dimensions<br>(includes air duct) | Height – 1958 mm  Width – 915 mm  Depth – 915 mm<br>Weight – 750 Kg | Design allows “brickwall” layout for better floor space utilization |


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Product Positioning – Ultrascale/Custom Space

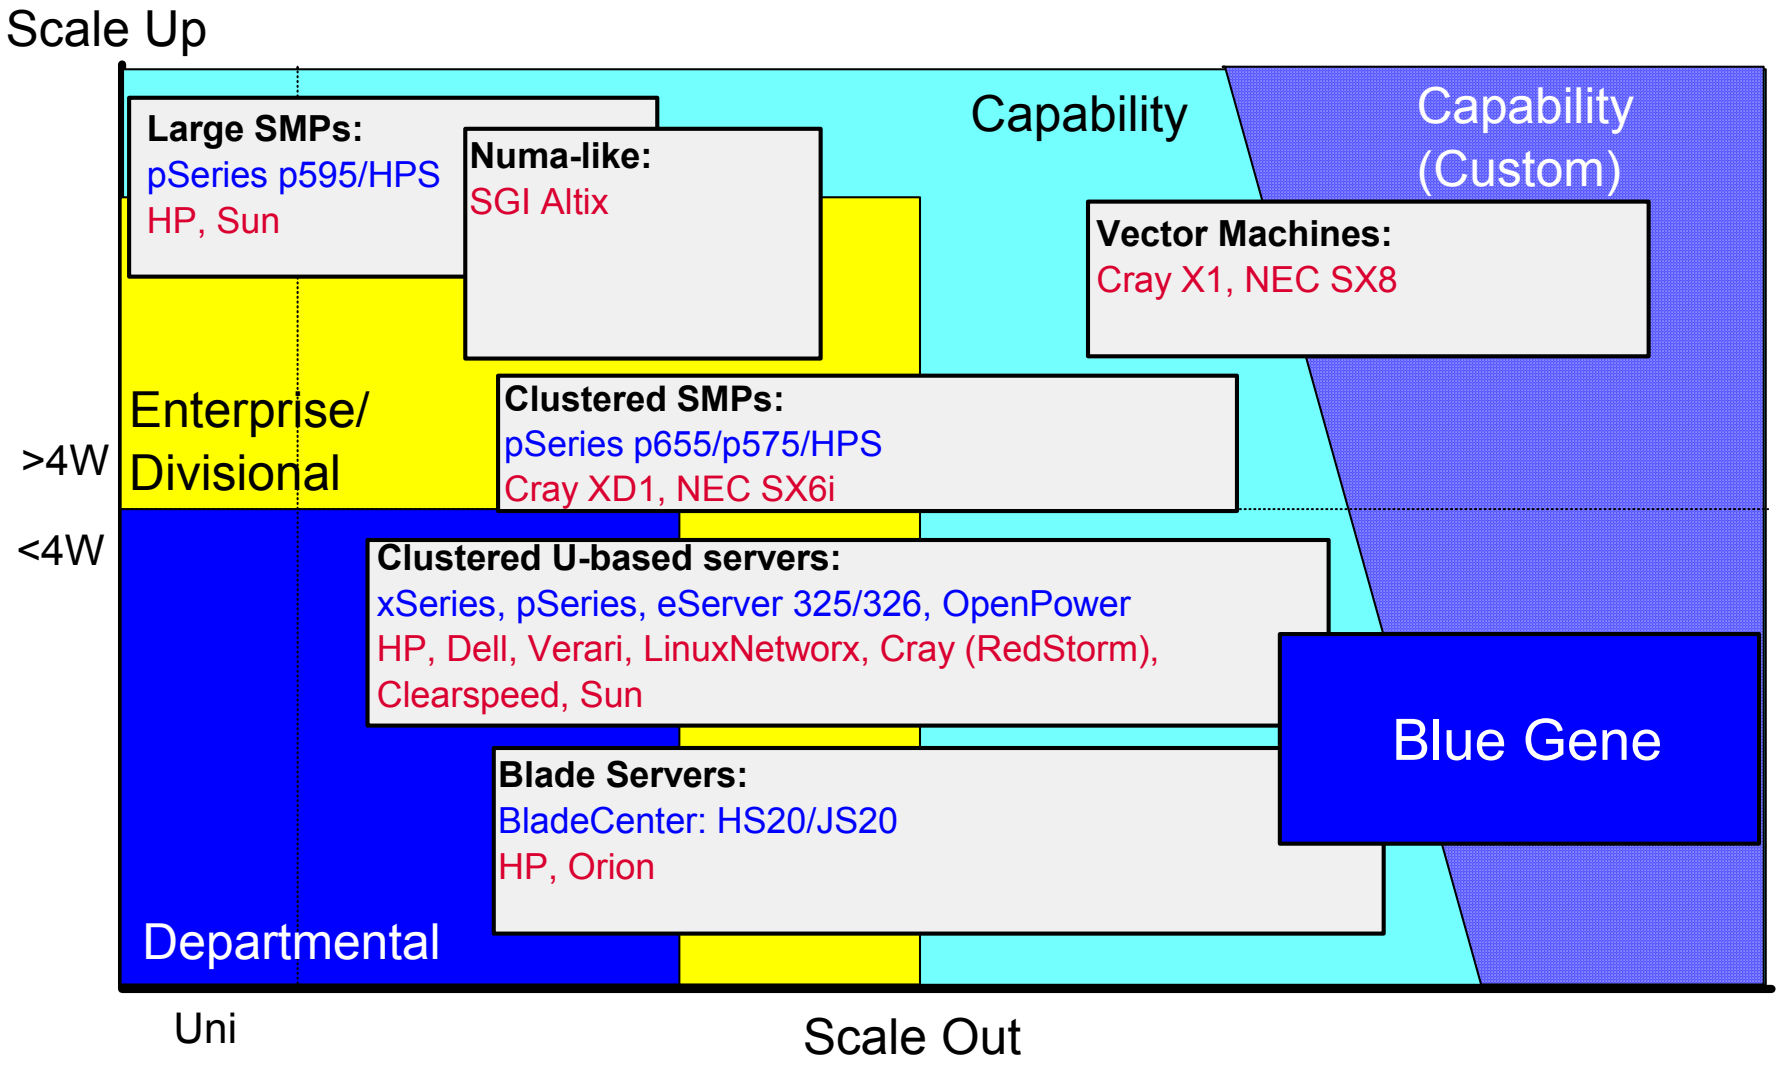



This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Making Blue Gene Accessible

- **Blue Gene On Demand**
  - IBM plans to provide access to Blue Gene through the Deep Computing Capacity on Demand (DCCoD) center
  - Clients with constrained budgets and limited requirements for access would be able to pay just for the amount of capacity reserved
- **IBM Global Financing support**
  - Leasing
  - Financing purchases
- **Reduced-Sized, Entry-Priced options**
  - Available to Blue Gene Consortium members on a limited time basis
  - 128-nodes in a partially populated rack
  - Upgradeable to 512 nodes

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Blue Gene – Product Directions

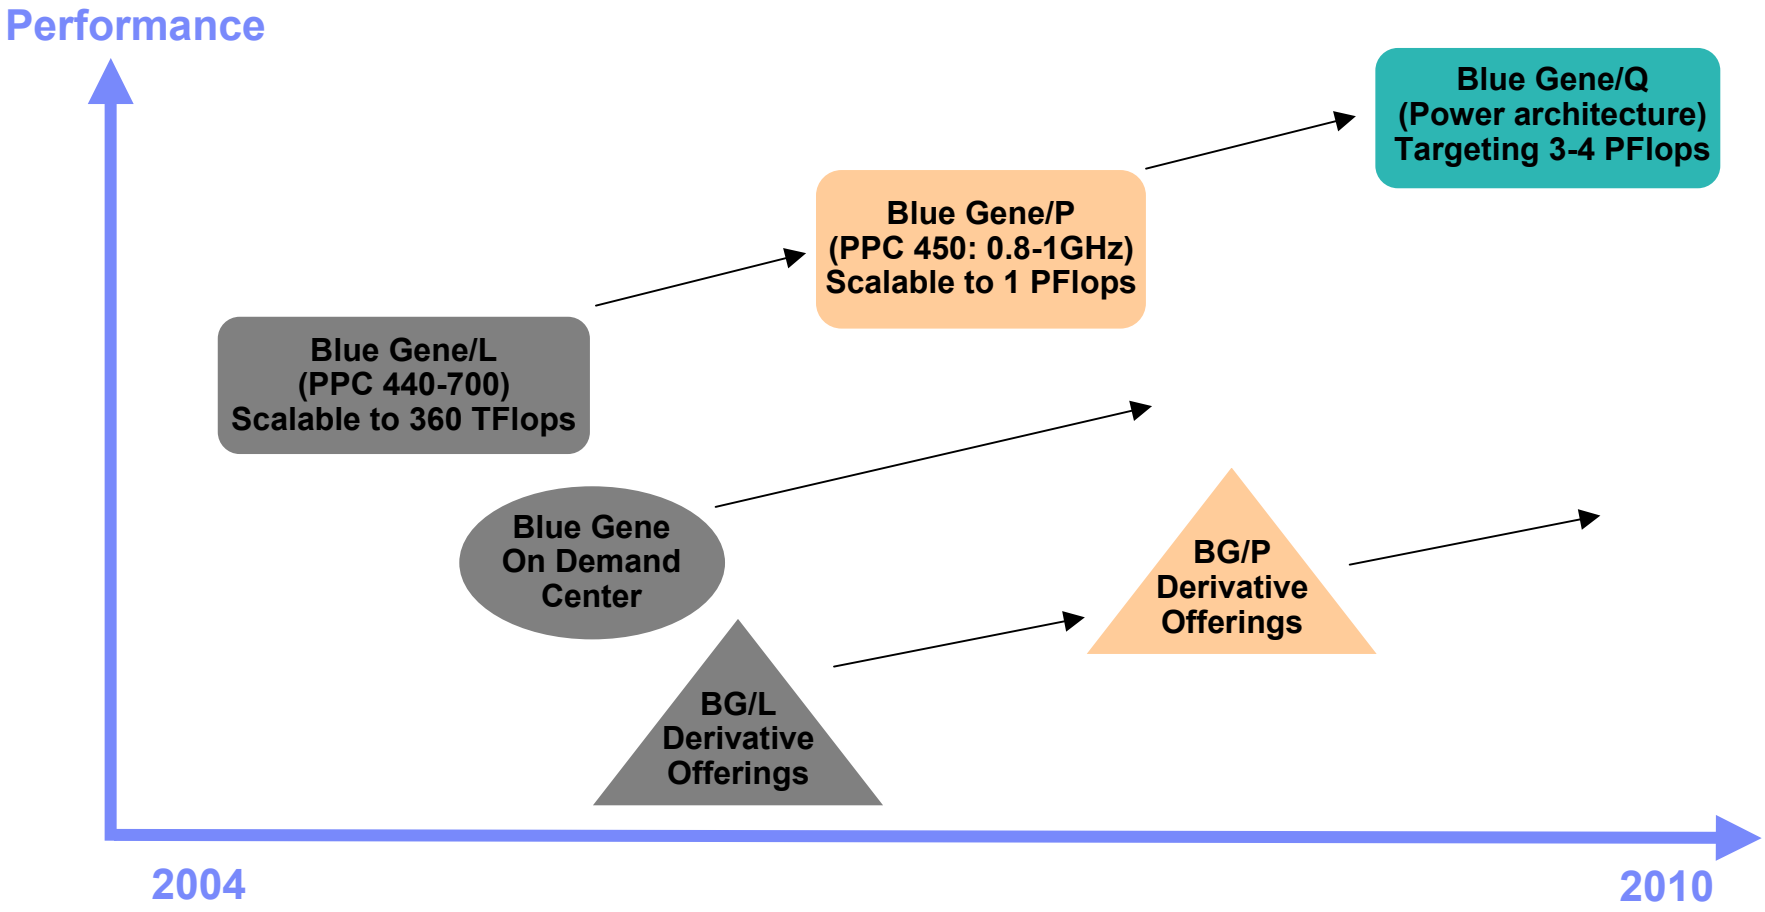



This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Closing Points

- **Blue Gene represents an innovative way to scale to multi-teraflops capability**
  - Massive scalability
  - Efficient packaging for low power and floor space consumption
- **Unique in the market for its balance between massive scale-out capacity and preservation of familiar user/administrator environments**
  - Better than COTS clusters by virtue of density, scalability and innovative interconnect design
  - Better than vector-based supercomputers by virtue of adherence to Linux and MPI standards
- **Blue Gene is applicable to a wide range of Deep Computing workloads**
- **Programs are underway to ensure Blue Gene technology is accessible to a broad range of researchers and technologists**
- **Based on PowerPC, Blue Gene leverages and advances core IBM technology**
- **Blue Gene R&D continues so as to ensure the program stays vital**

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# BlueGene Pipeline/Leads by Segment by Geo

| | ▪AG | ▪AP | ▪EMEA |
|---|---|---|---|
| **Govt/Research & Defense** | ▪LLNL<br>▪Argonne National Lab<br>▪Los Alamos National Lab | ▪AIST | ▪ASTRON |
| | | ▪KEK High Energy Accelerator Research Organization<br>▪Korea Instutite for Science & Technology (KIST) | ▪Atomic Weapons Establishment (UK)<br>▪ECMWF (Euro Cntr for Medium Range Wx Forecasting)<br>▪LOIS (LOFAR Outrigger in Scandinavia)<br>▪NIF / Hungary (National Info Infrast Dev Pgm) |
| **Higher Ed/ Research** | ▪SDSC<br>▪Boston University<br>▪NCAR | | ▪U of Edinburgh<br>▪EPFL<br>▪FZ/Juelich |
| | ▪ACEnet (Atlantic Canadian U's)<br>▪Iowa Universities / Ames Research<br>▪Maui HPCC<br>▪Purdue U<br>▪Sharcnet (South Central Ontario U's/Canada)<br>▪HPCVL (Queen's University/Canada)<br>▪StonyBrook U<br>▪U of Louisiana<br>▪Uof Minnesota | ▪Seoul National University (SNU) | ▪HECToR (UK Supercomputer)<br>▪Royal Institute of Technology (Stockholm)<br>▪Trinity College / TCHPC (Ireland) |
| **Life Sciences/ Pharma** | | | ▪San Rafaele Med Cntr |
| | | | ▪CBS/Denmark Center of Biological Sequence Analysis |
| **Commercial (Non Public Sector)** | | ▪NIWS | |
| | ▪Raytheon | | ▪CGG<br>▪DaimlerChrysler (DCX GmbH)<br>▪Mercedes-McLaren (Germany) |

***Note: Some leads not yet validated***



This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Blue Gene - Call to Action

- **Identify opportunities in selected segments**
- **Contact E&TS Principal(s) and Deep Computing team**
- **Team to validate leads and prioritize opportunities**
- **Present BlueGene proposal to customer**
- **Clarify with customer:**
  - Applications
  - Requirements
  - Funding Amount, Availability and Process
- **Team to close validated opportunities**

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Cell Processor Development
# Sony, Toshiba and IBM's Microprocessor Alliance

- **March, 2001: "Cell" Processor Development Alliance Formed**
  - Shared investment of more than $400 million in the next five years to design a "supercomputer-on-a-chip."
  - Joint research and development for an advanced chip architecture to support a new wave of devices in the emerging broadband era.
  - "STI" Design Center is located in  Austin Texas
- **Alliance combines:**
  - Sony/SCEI's vision and strong leadership in the computer entertainment world
  - IBM's unparalleled computer and semiconductor technologies
  - Toshiba's extensive capabilities in system LSI (large-scale integration), especially for consumer applications, the companies

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Cell Processor Characteristics

- **Vast floating point capability**
- **Massive data bandwidth**
- **Parallel processing architecture**

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Sony, SCE and IBM's System Collaboration

## *Cell Processor Based Workstation*

- Next Generation content development platform for the digital media industries
- Scalable, supercomputer-like performance
- Common Software Development Environment for digital animation and video game industries
- *First* Cell Processor based system product !

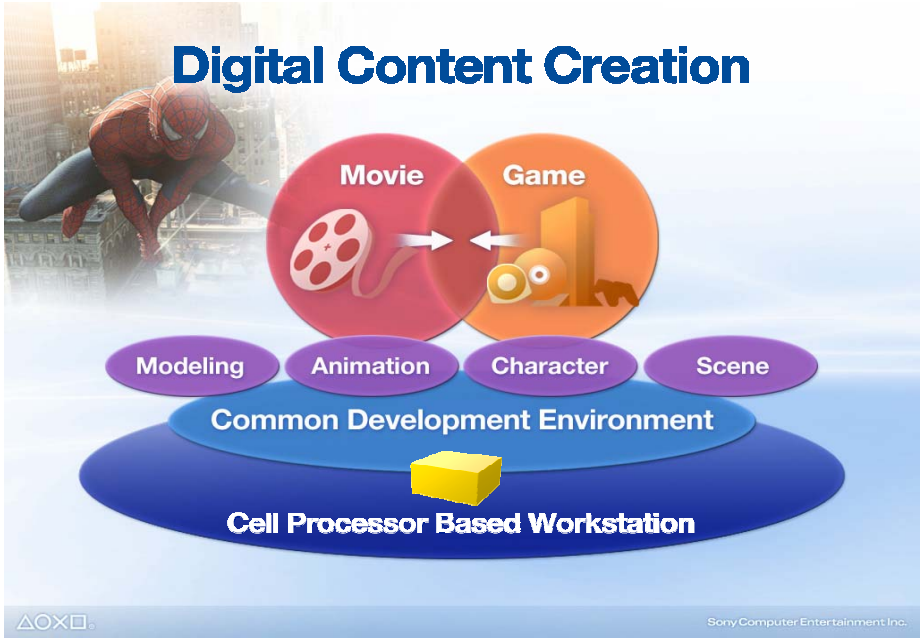


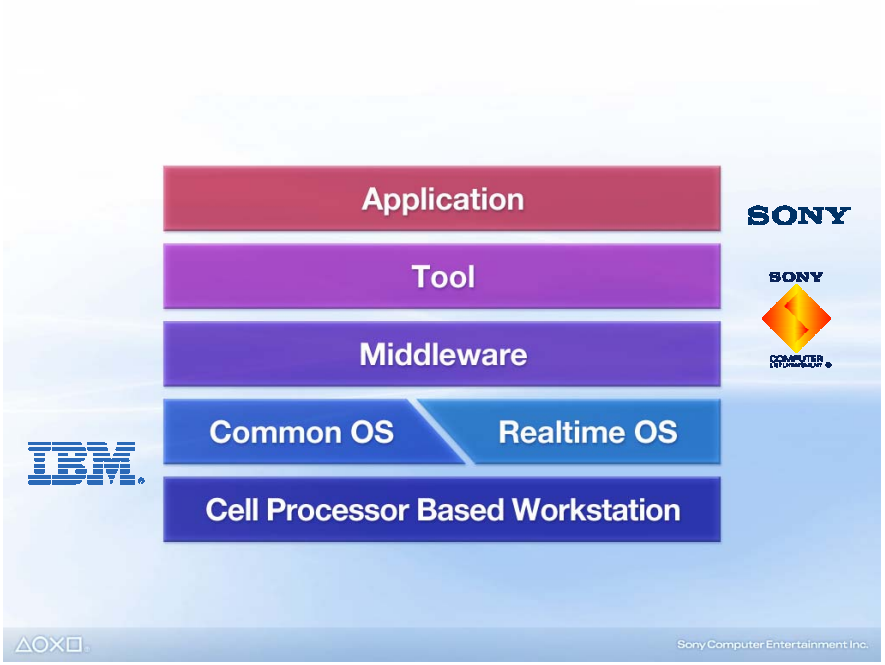


***Announced May 11, 2004***

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Cell Processor Based Workstation Overview

***•High Performance Digital Content Creation Workstation***
***•Key Element of the IBM-Sony/SCE Integrated Digital Content Creation Environment***

- **Blade**
  - Dual Cell Processors (SMP), Support Logic, Memory, Storage
  - PCI Express 4X option port
  - BladeCenter Interface (Based on IBM JS20)
  - Infiniband 4x interconnect
- **Chassis**
  - Standard IBM BladeCenter form factor with:
    - 7 Blades (for 2 slots each)
    - 2 internal switches (1Gb Ethernet) with 4 external ports each
  - Separate, external Infiniband Switch with optional FC ports
- **Software**
  - Cell-aware
    - Linux OS
    - Tool chain and compilation support
  - Additional IBM software development tools and cluster management software

**Preliminary**

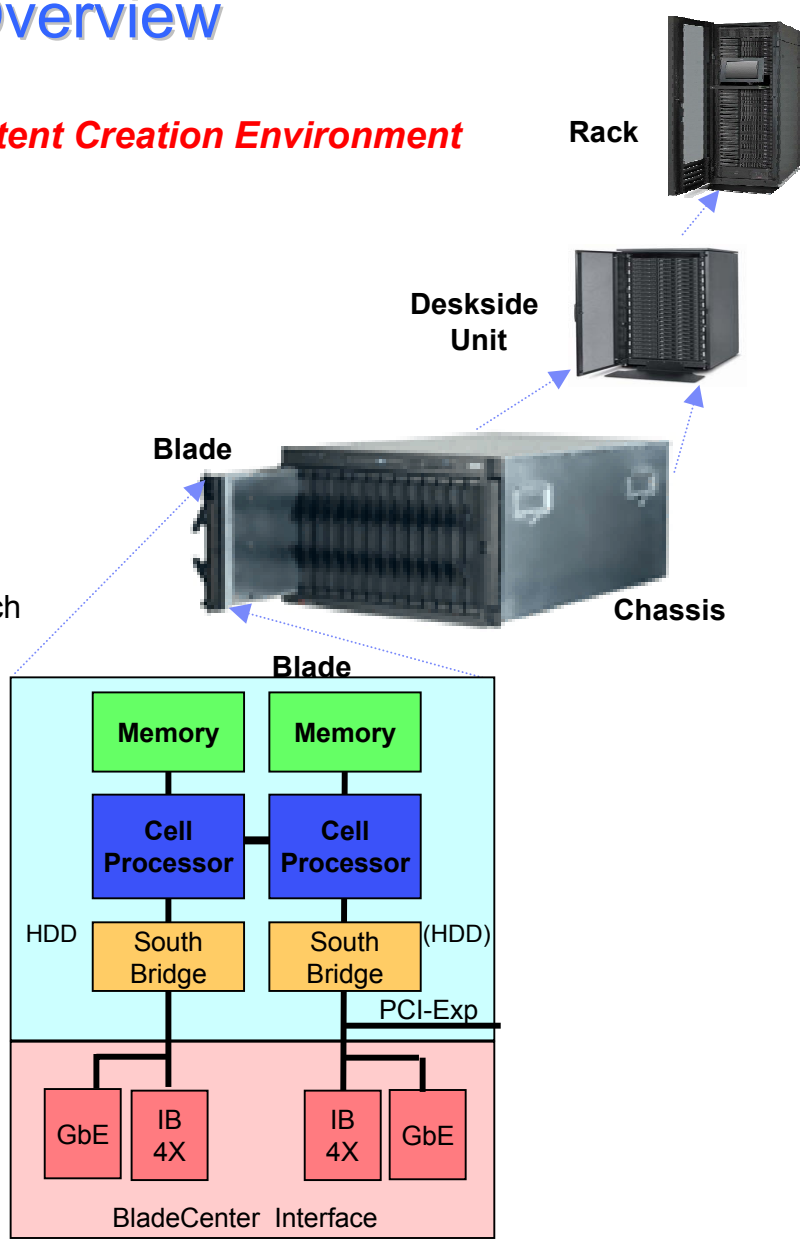



This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Strategic Directions for IBM's Cell Processor-Based Systems

**Workstations**

**Servers**

**Embedded Processors**

**Real Time Applications**

***→ An Open, Scalable architecture that will be a platform for many of today's and tomorrow's ideas!***

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Deep Computing Visualization

## *Transforming data into insight*

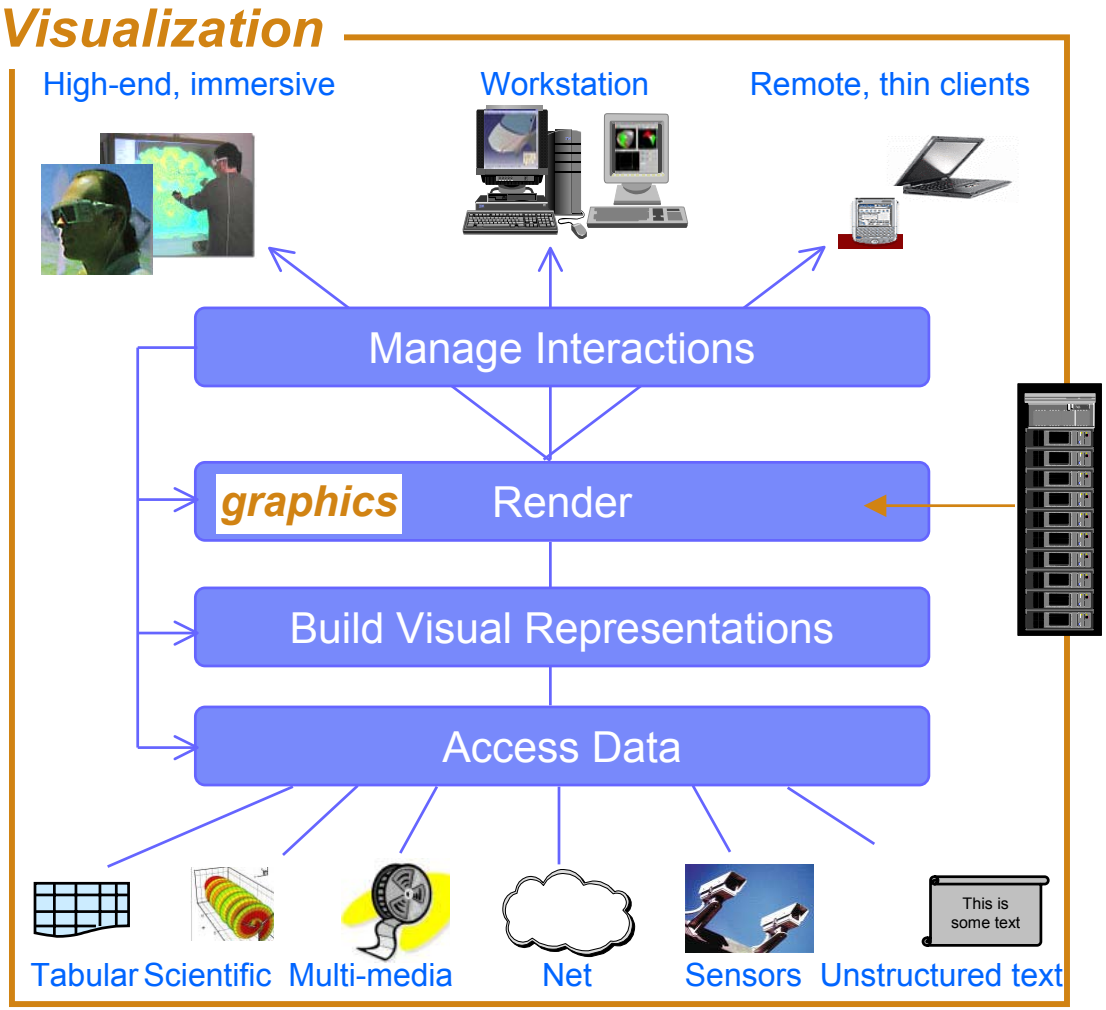


- ***DCV*** provides graphics for high end and remote applications
- Optimal price/performance
- Leverages cluster trends and aggressive gains of commodity graphics and networks

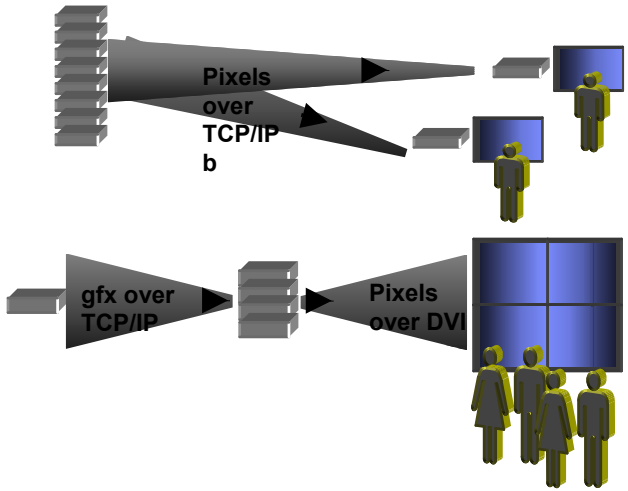


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Why is visualization fundamental to IBM ?

- **DCV market opportunity >$50M for 2005 (and >$500M market potential)**
  - Visualization is a catalyst for company growth
- **Eliminates NO-Bids and targets specific industries**
- **Extends the infrastructure and data management plays end-to-end**
  - Visualization enables product and services drag
- **Enables the sales of IBM products and services via new paradigms**
  - Visualization supports the On Demand initiative
  - Remote visualization and collaboration is a new workflow paradigm

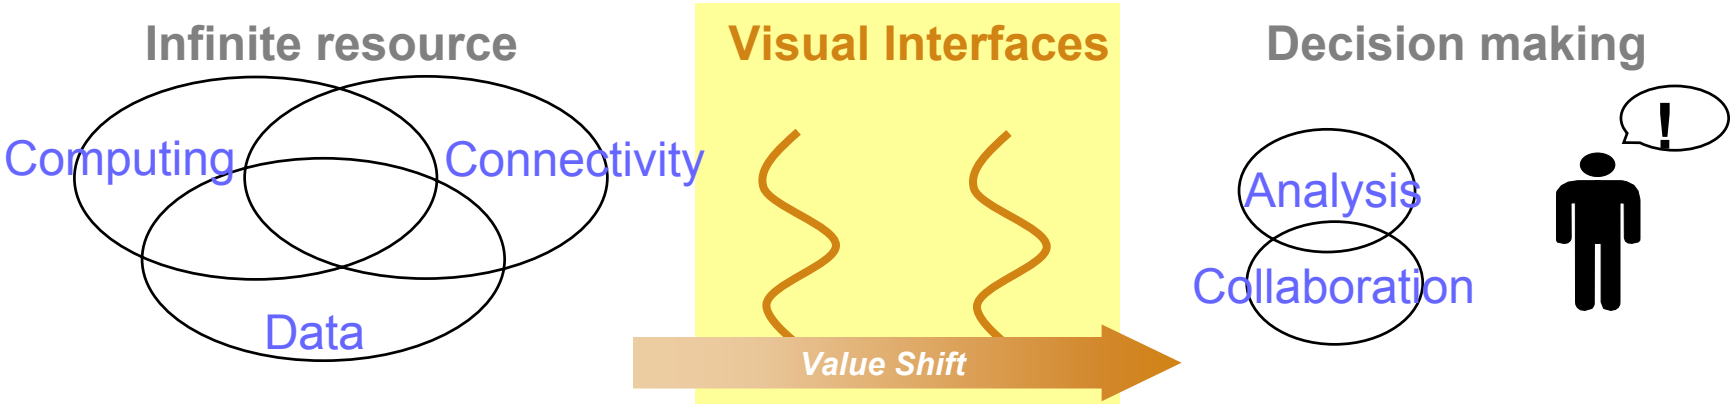


## ….Transformational Insight, Not Transactions


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# The Role of Visualization in Deep Computing

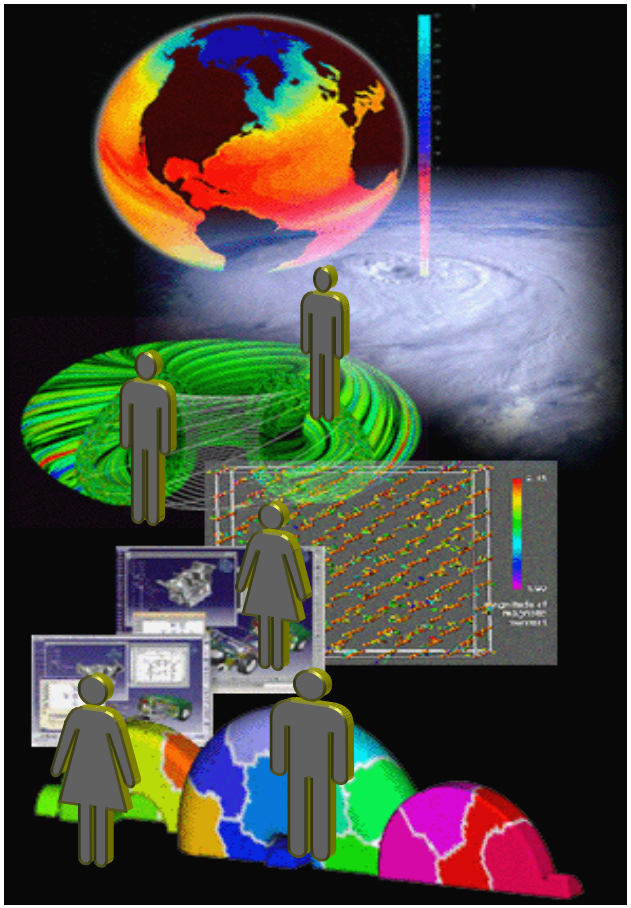

- Navigate and interface data
  - Data exploration
  - Computational steering
- Transform digital data into analog information
  - Overall data reduction
- Couple computing to human experts
  - Enable decision support
  - 50% of human neurons involved in vision
- Enable collaboration
  - Bring visualization to desktops around the globe


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Scalable Visual Networking: SVN

## Immersive Visualization

- Run unmodified application on master workstation and transport graphics streams to display-attached rendering nodes
- Multiple graphics transport options offer performance alternatives
- Scale resolution of OpenGL windows for use with rich-pixel, multi-pipe environments:
- Stereo capable for stereo enabled applications.

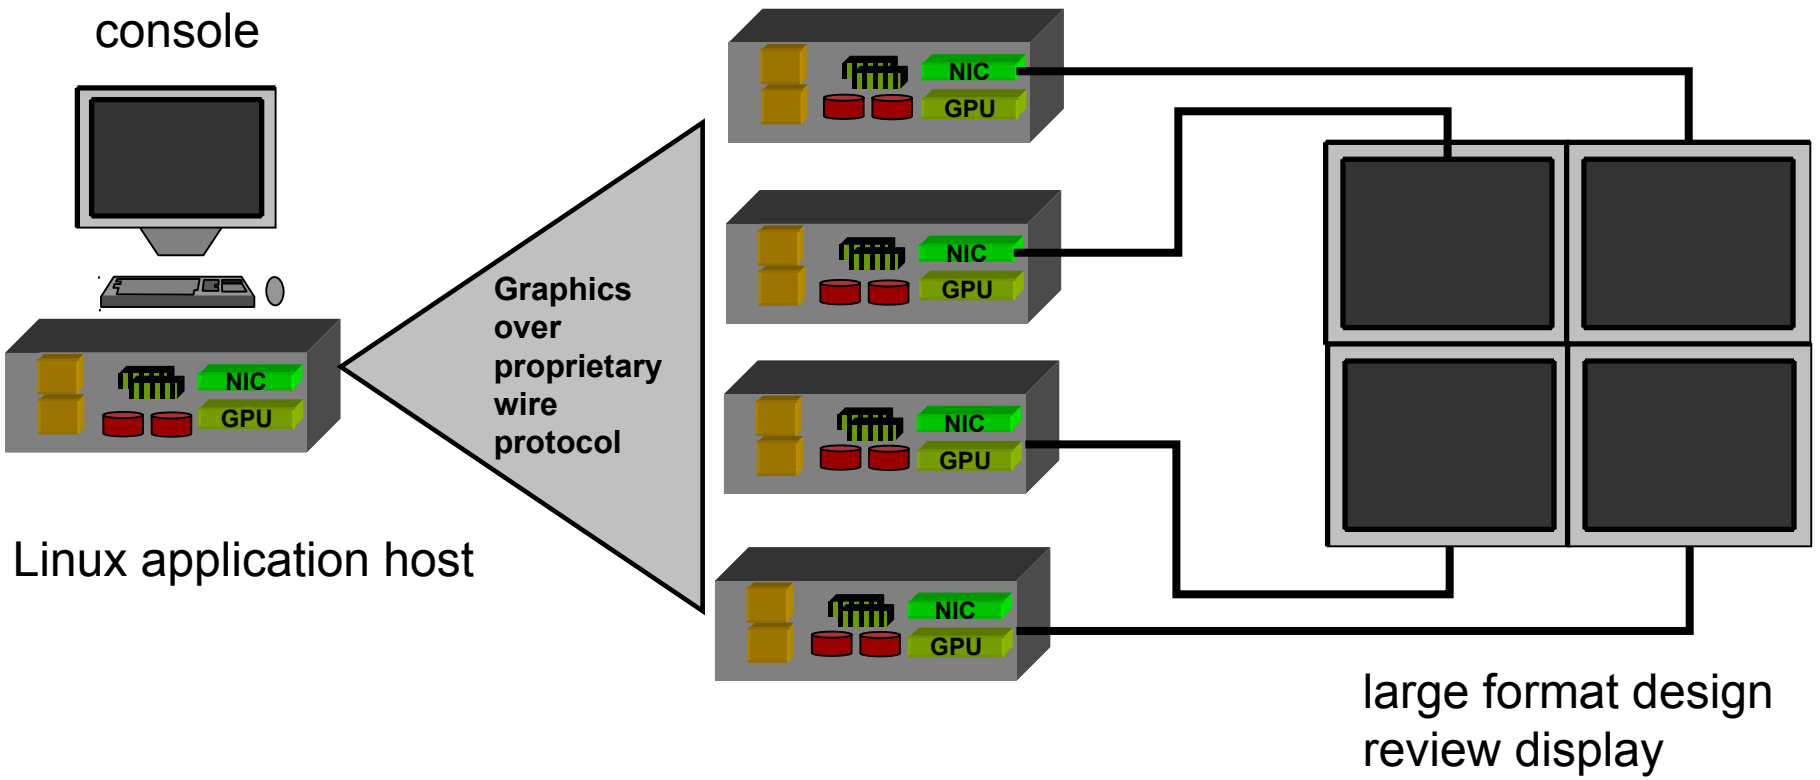



This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Remote Visual Networking: RVN

## Long-haul application delivery over networks

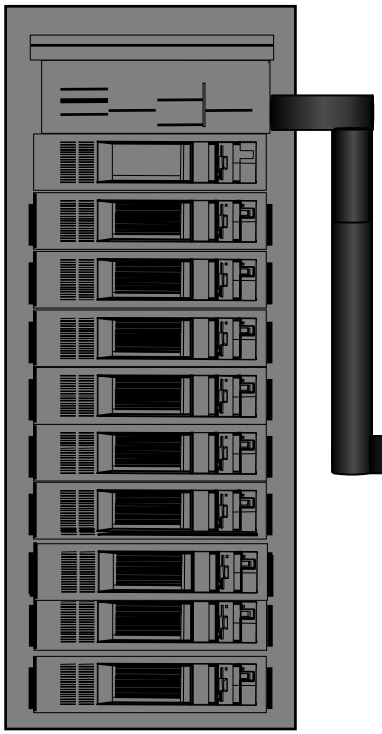

- Utilize end-station for access to a remotely executed and rendered OpenGL application
- Enables re-location or consolidation of Linux workstations

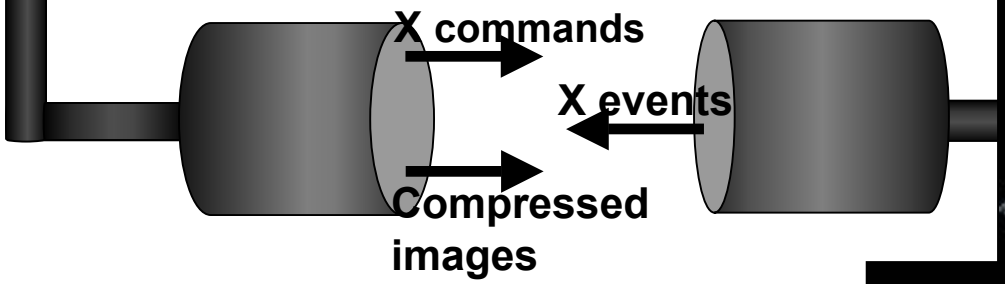


- Can be used for multi-user collaboration of a single application session
- Collaborators can join and detach from a collaborative session arbitrarily

DCV cluster managed like a pool of workstation resources

ThinkPad with both installed applications and remote applications



This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Inside The Deep Computing Visualization Cluster

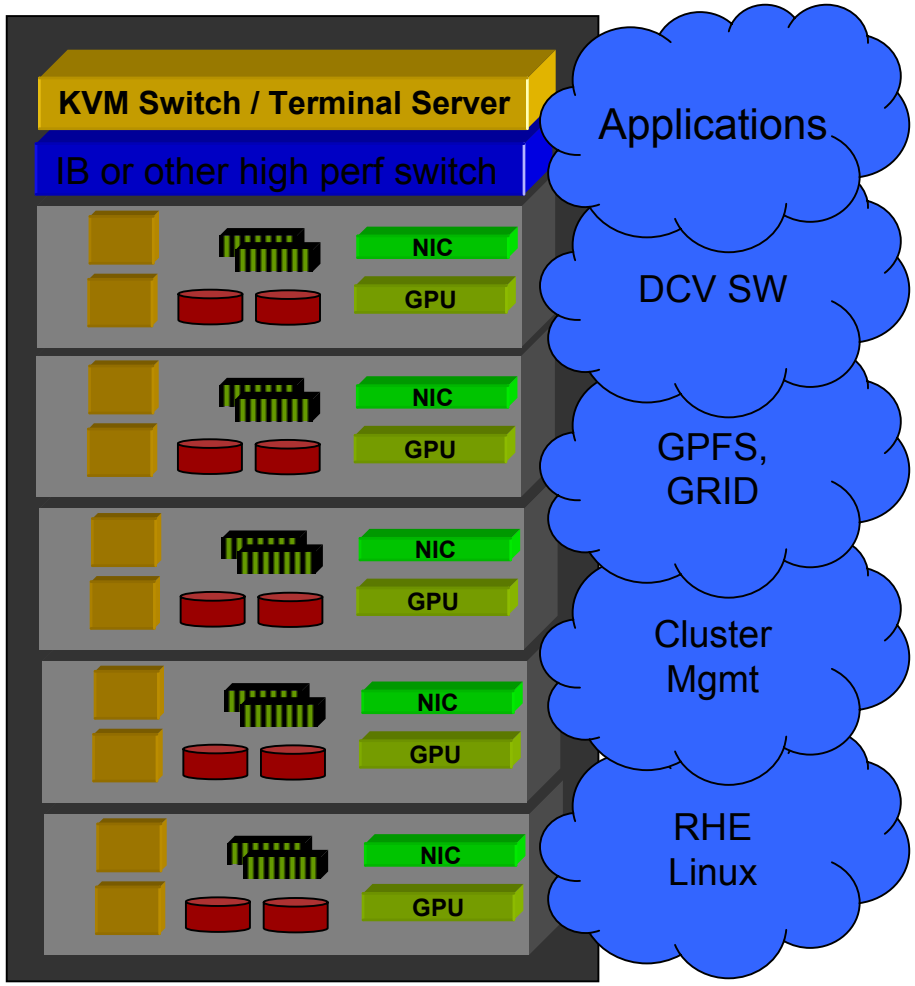


- Rack mounted IBM IntelliStation A Pro or Z Pro workstations serve as rendering/media nodes
- AGP 8X/PCI-Express Graphics for commodity hw accelerated rendering
- InfiniBand (IB) or Gigabit Ethernet (GE) as cluster interconnect to transport geometry, synchronization, events, and pixels
- Multi-user, multi-application aware cluster
- Processors available for batch computation based on availability
- Can be scaled or integrated into other clusters


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Target Markets for DCV

| Sector | Target |
|---|---|
| **Industrial** | **Petroleum Exploration & Production (C&P)**<br>**Product Lifecycle Management (PLM)**<br>**Product Data Management (PDM)**<br>**Computer Aided Engineering (CAE)** |
| **Public** | **Government Research & Classified**<br>**Higher Education Research**<br>**Weather/Environmental Sciences** |
| **Life Sciences** | **Research & Development**<br>**Medical Imaging** |
| **Digital Media (explore)** | **Digital Content Creation, Management & Distribution** |


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# What do the Analysts think about DCV?

- **"IBM is targeting three types of potential customers with DCV:**
  - For users of traditional proprietary graphics systems…"
  - For users of personal workstations and PC's, DCV aims to provide a more scalable solution…"
  - For users who wish to incorporate visualization into their workflows for the first time…"

- **"IBM's greatest opportunity lies in the possibility that the industry could shift to this new architectural paradigm [scalable graphics] as the preferred technique for advanced visualization, with IBM poised to take advantage of the new dynamic."**

*Addison Snell and Christopher G. Willard, Ph.D, IDC, January 2005*


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# DCV Ordering and Pricing

- **Via Passport Advantage**
- **List price (NTE): US$2500 / PWS processor**
- **Product ID (PID): 5724-K69**
- **Warranty: 1 Year (minimum)**
- **Software Maintenance: 1 Year & Renewable**
- **Planned announce date is February 15, 2005**

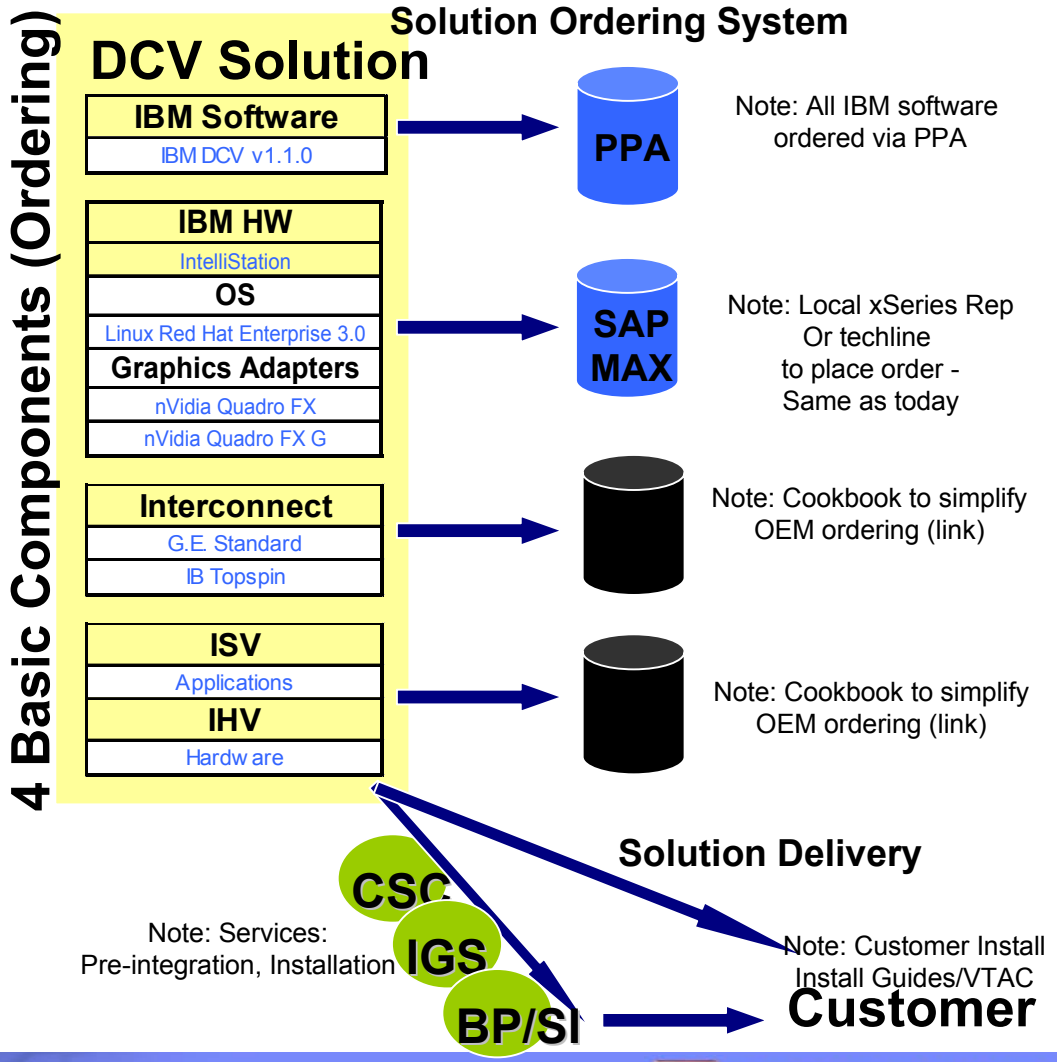



This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Deep Computing Visualization Summary

- DCV is a scalable, flexible, modular, readily deployable visualization infrastructure based on commodity hardware and differentiating software
- DCV uses IntelliStations, IB or GE networks, commodity graphics, and virtualization software
- DCV has a capability called Remote Visual Networking or RVN that enables remote collaboration and workload consolidation with applications on thin clients
- DCV has an immersive, high-end capability called Scalable Visual Networking or SVN that enables CAVES, Powerwalls, DMUs, and data centers

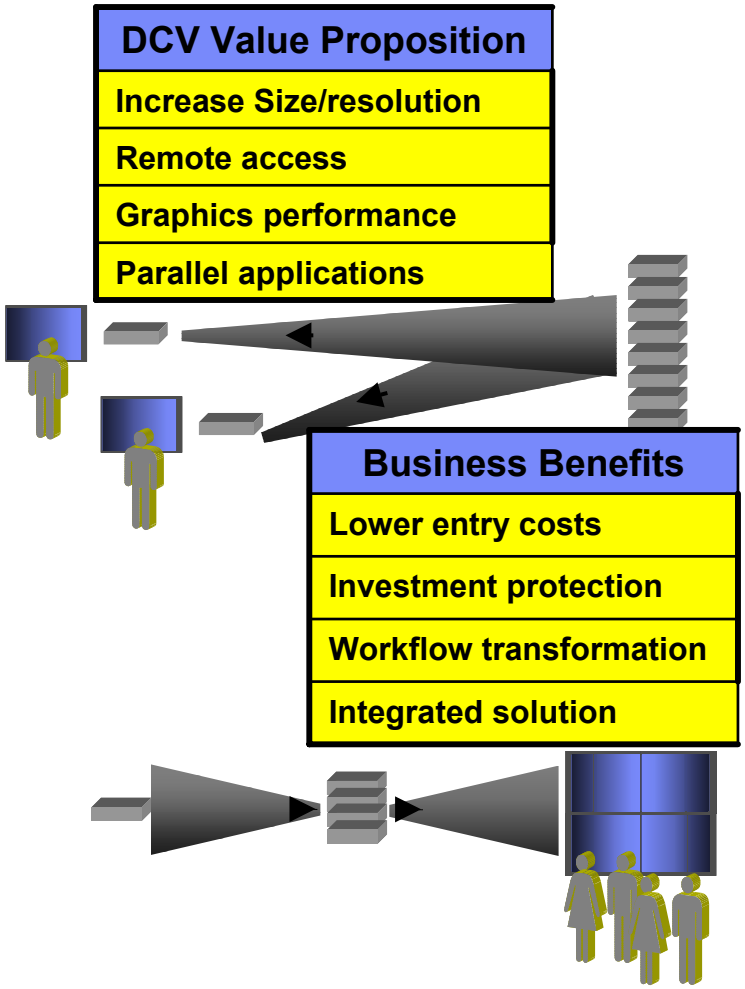


Visualization is core to innovation and insight…


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Deep Computing Capacity on Demand

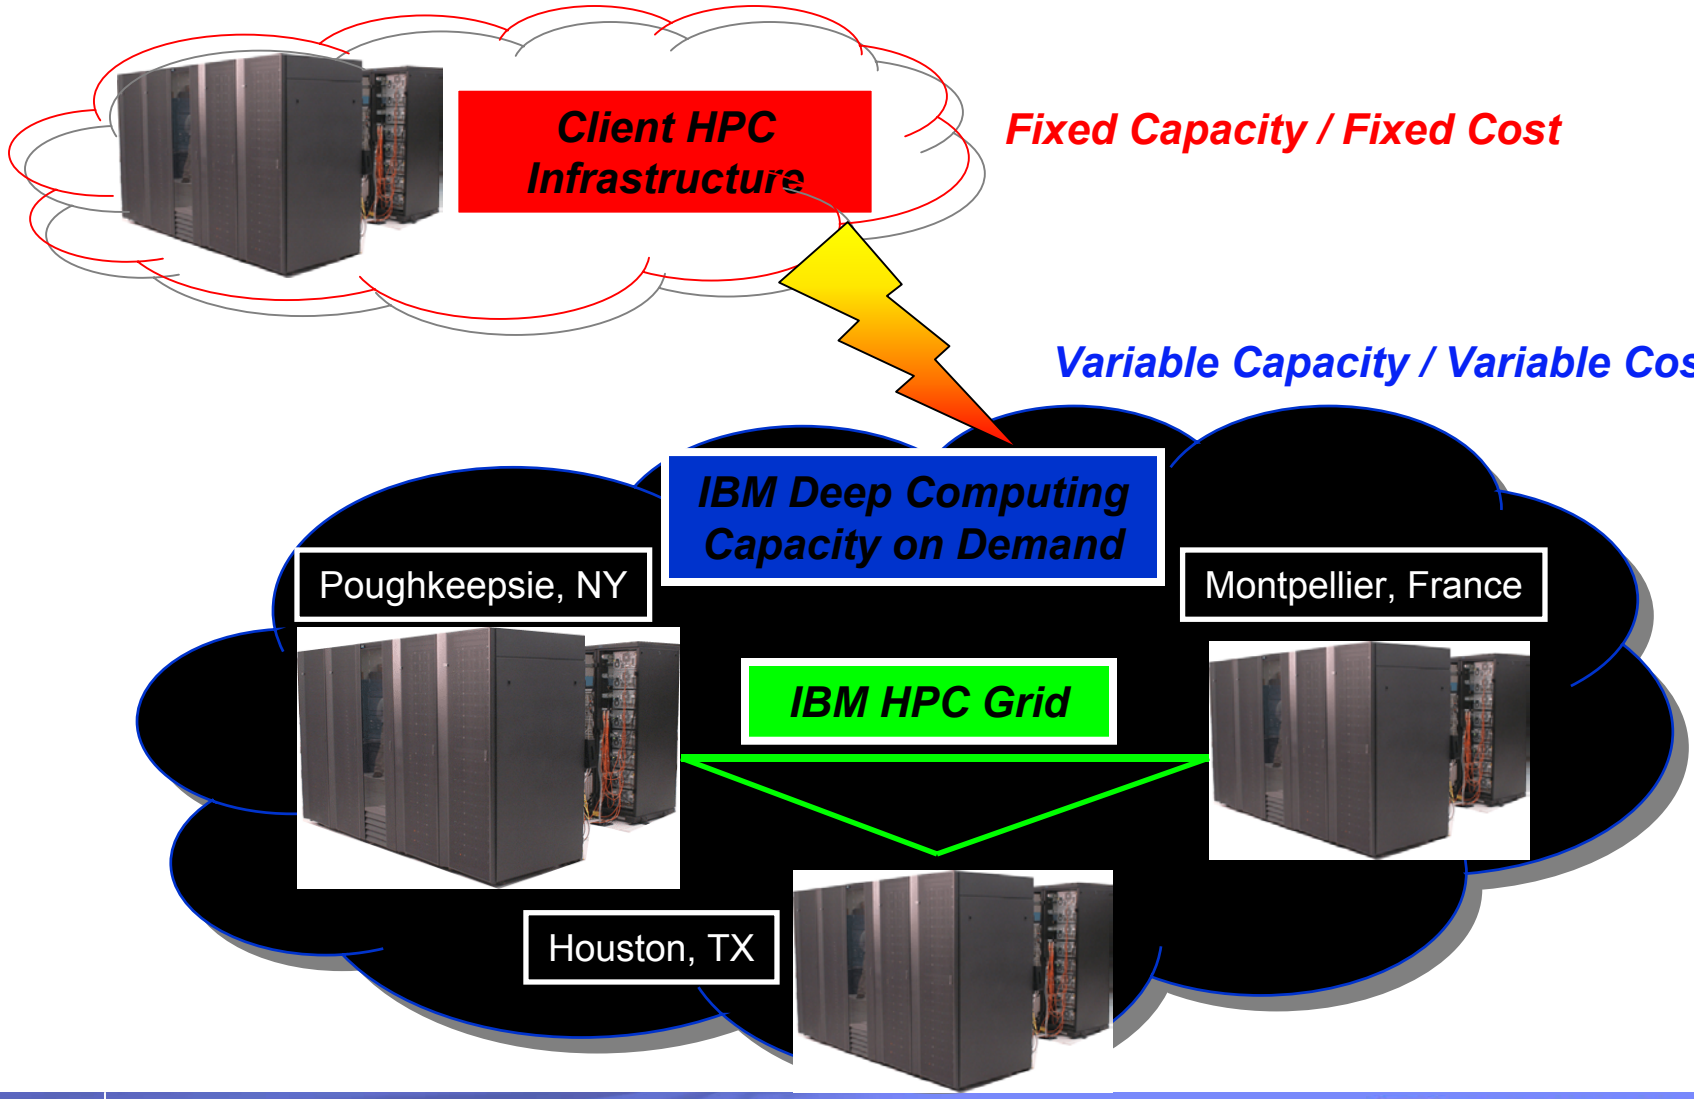


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Flexible Financial and Delivery Models

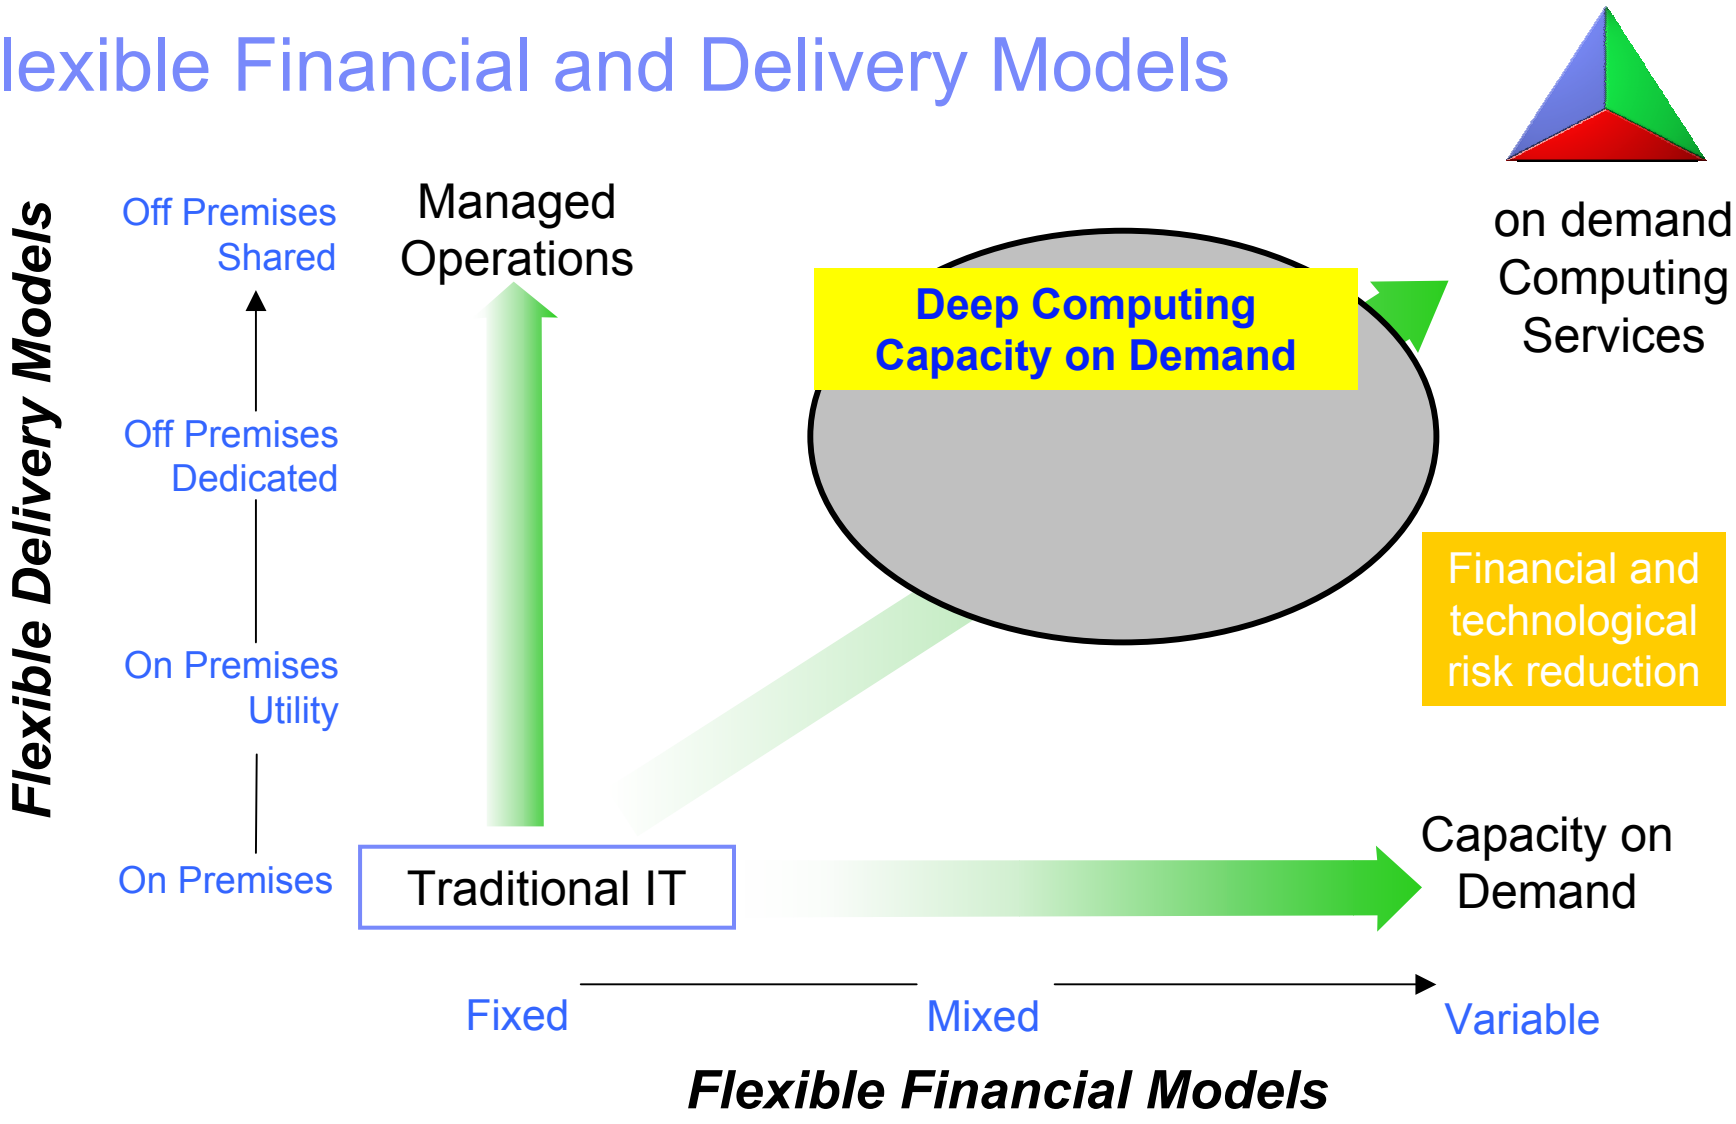


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Customer Value

## Scalable Capacity

Access "unlimited" HPC capacity to

- ☑ Compete at a scale that may not have been possible previously
- ☑ Solve larger, more complex problems
- ☑ Perform more detailed analysis and more iterations in the same amount of time
- ☑ Finish projects in less time

## Time-to-Value

- ☑ Respond faster to new business opportunities
- ☑ Respond faster to keep-the-business-running peak workload demands
- ☑ Improve deployment time for incremental HPC infrastructure and capacity
- ☑ Improve time to productivity and return on investment

## Variable Costs

- ☑ Avoid upfront capital acquisition and long term fixed IT cost lock-in
- ☑ Better align IT costs with actual consumption
- ☑ Deliver immediate and predictable ROI
- ☑ Simplify IT chargeback to projects and business units

## Risk Mitigation

- ☑ Avoid over-capacity and under-utilization
- ☑ Avoid under-capacity and inability to respond
- ☑ Avoid financial and management responsibilities of owning a supercomputer
- ☑ Avoid acquisition, deployment, upgrade and maintenance cycles
- ☑ Avoid technology obsolescence

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Target Markets & Workloads

| Industry | Target Workloads |
|---|---|
| Petroleum | Seismic analysis, reservoir modeling/simulation |
| Electronics | Design verfication, auto test pattern generation, tape-out, mask generation, routing, design rule checking, optical proximity correction (OPC) |
| Digital Media | Animation, online gaming |
| Life Sciences | Drug discovery, bioiformatics, biotechnology, proteomics, genomics, clinical trials |
| Financial Services | Risk management, compliance, portfolio/wealth management |
| Automotive | CAE, computational fluid dynamics, crash analysis, structural analysis |
| Aerospace & Defense | CAE |
| Government Research | Life Sciences, weather, energy/nuclear |
| Higher Ed. Research | Life Sciences |

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# IBM Offerings for Deep Computing Capacity on Demand

## CORE INFRASTRUCTURE

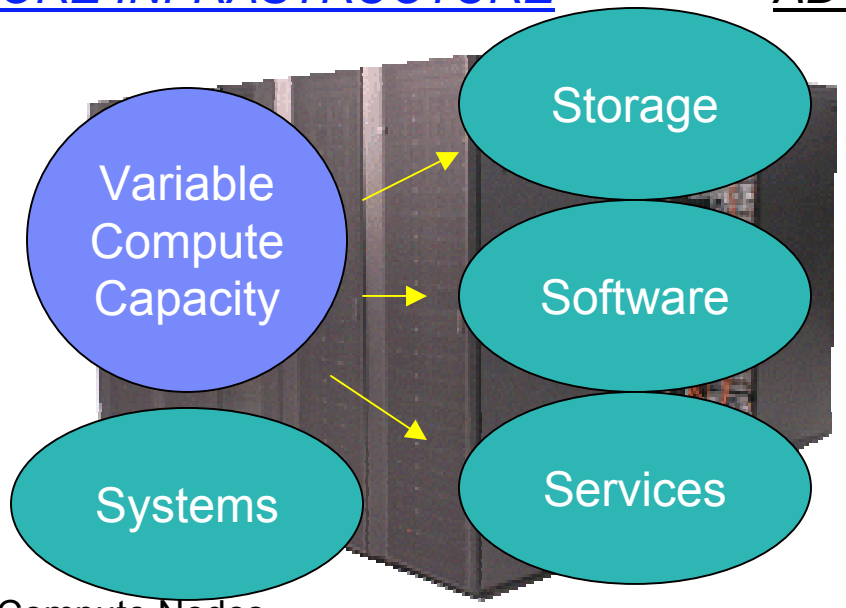


- Compute Nodes
  - IBM @server™ Cluster 1350
    - xSeries® Intel® Xeon™ 32-bit technology
    - AMD Opteron™ processor 32/64-bit technology
  - IBM pSeries® POWER4+™ 64-bit technology
- Interconnect: 10/100 & Gigabit Ethernet
- Management Node
- Cluster Systems Management (CSM)
- Virtual Private Network (VPN)

## ADVANCED FEATURES AND SERVICES

- IBM @server Storage Nodes
- IBM TotalStorage® Fiber Channel and SCSI Disk

- IBM General Parallel File System (GPFS)
- Job scheduling and workload management (*)
- Grid-enablement capability

(*) Planned Availability

- Advanced monitoring & management
- Software customization & deployment
- Storage/data management (e.g., backup)

## BASE SERVICES

- Facility
- Provisioning
- Deployment
- Virtual Private Network
- Security
- Monitoring
- Maintenance
- Help Desk
- IBM Representative

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Deep Computing Capacity on Demand - Capacity

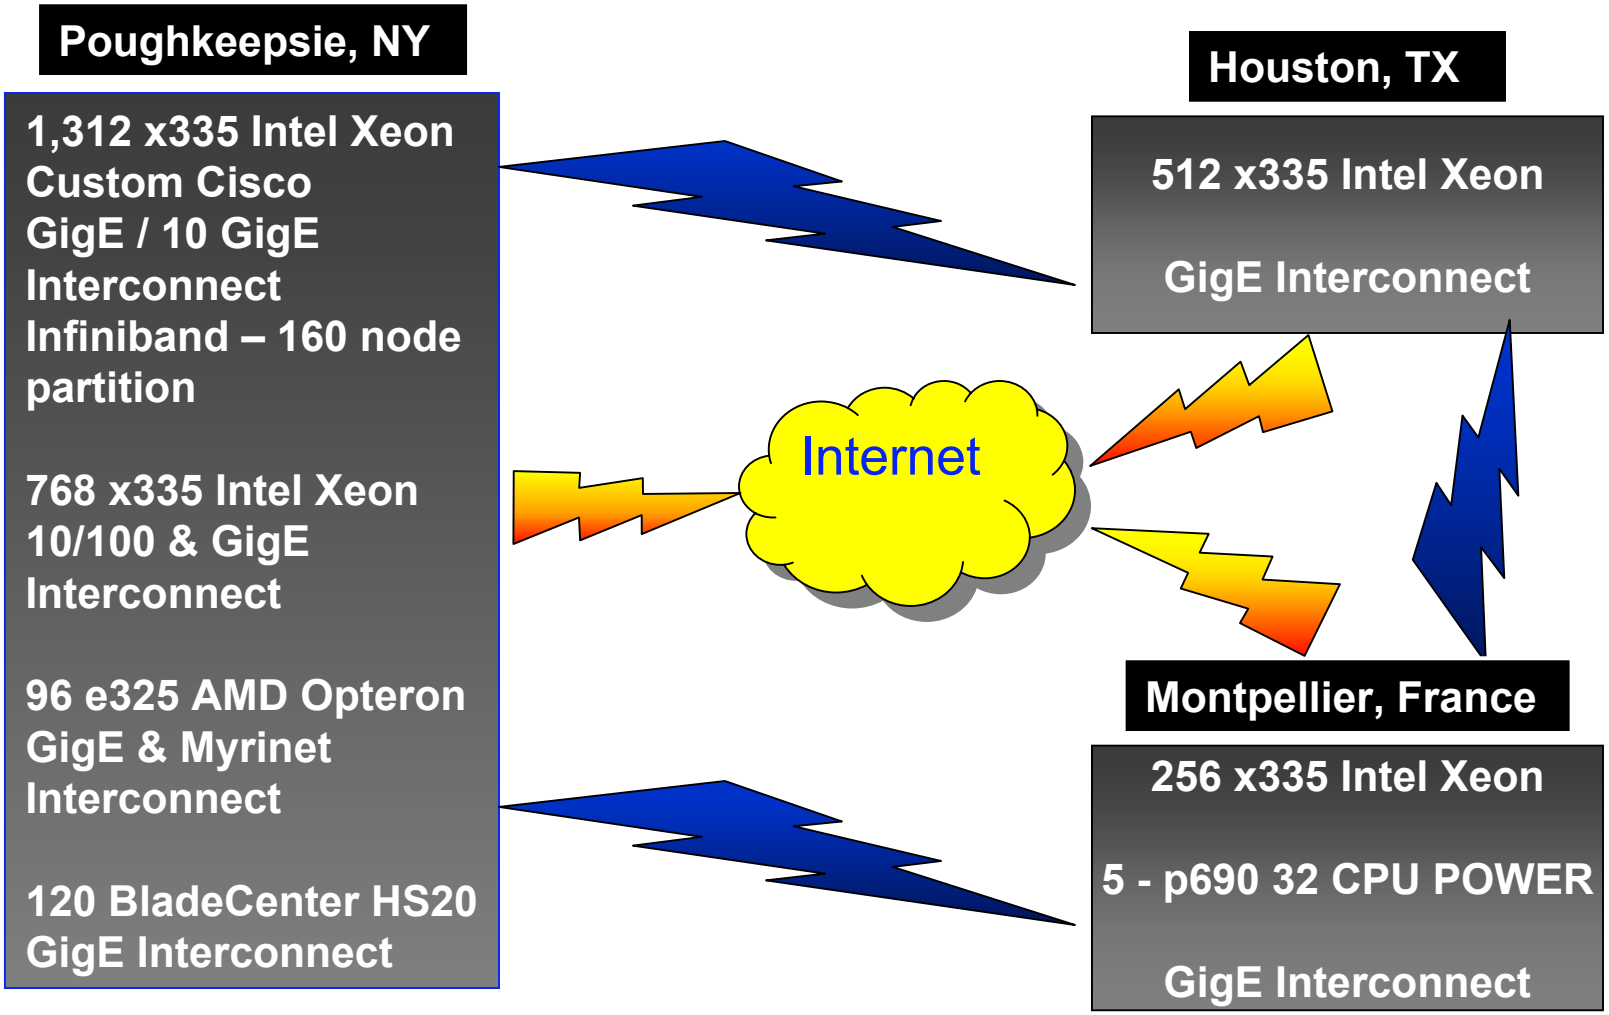


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Shared Resources with Virtual Cluster Isolation

- **Multiple customers access "virtual" clusters made up of compute and storage resources**
  - Allocations subject to availability
- **Customer controls virtual cluster**
  - Root access to system
  - Install and run software
  - Use CSM to manage system
  - Advanced features available
- **Highly secure and resilient infrastructure**
- **Customer components**
  - Software and licenses
  - Data
  - Hardware

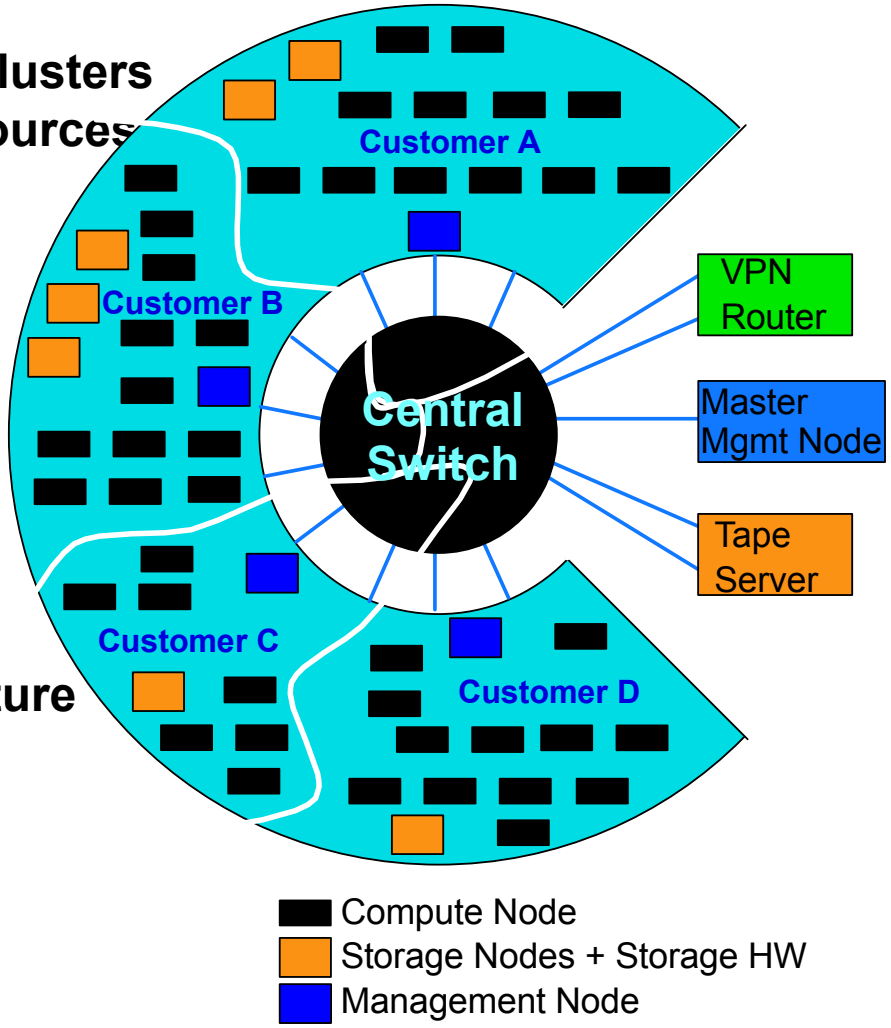


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Two-Phase Implementation

- **Base Presence & Access ("Membership")**
  - 4 nodes (2 compute, 1 storage, 1 management)
  - Optional dedicated storage
  - VPN connection
  - Annual Facility Charge

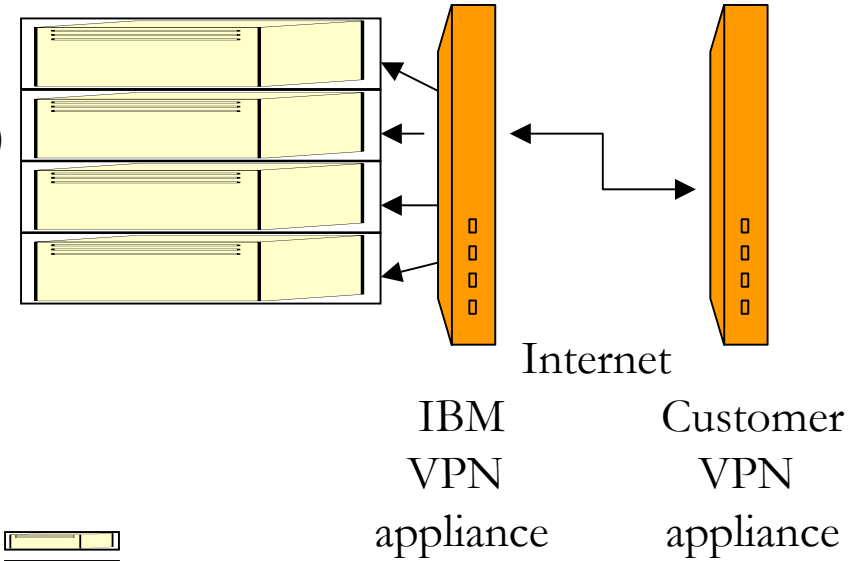


- **Engagement (Peak Demand Capacity)**
  - Groups of 16 nodes, storage, enhanced network connectivity
  - Weekly reservations
  - $/CPU-Hour metric
    - Based on volume, duration, advanced reservation, additional commitments
    - $/GB-Week for storage

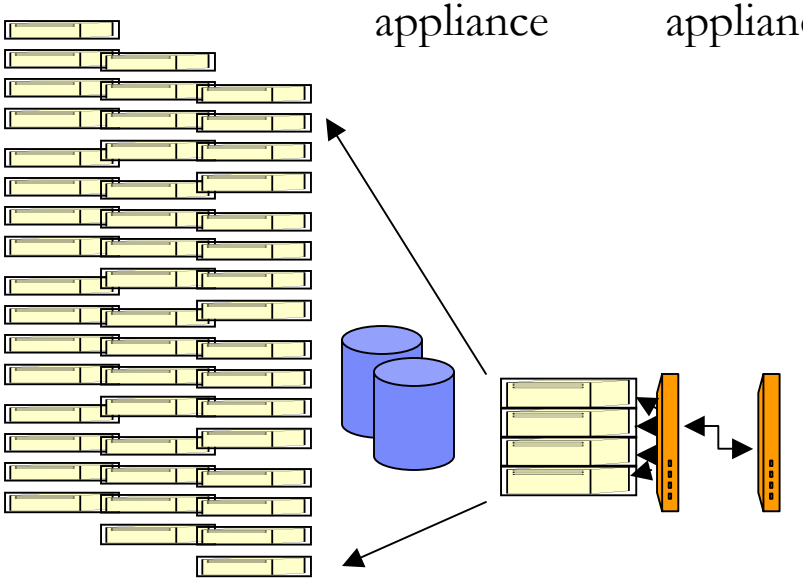

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Grid-Enabled Dynamic Scheduling & Provisioning

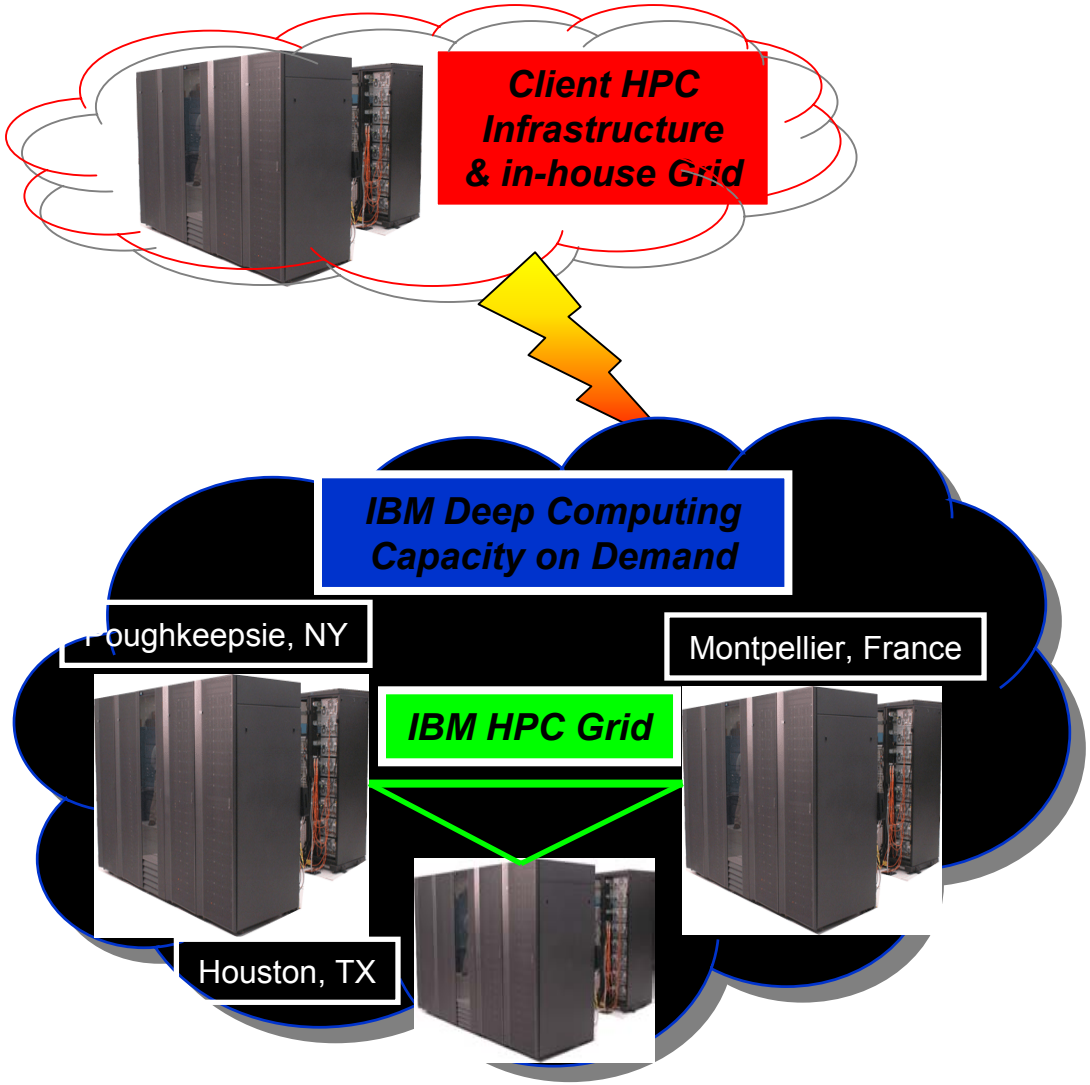


## Tap into IBM's HPC Grid

- ✓ Extension of client's in-house HPC infrastructure and Grid
- ✓ New 'utility-like' model offering
  - in addition to existing and continuing 'reservation' model
- ✓ Increased flexibility in choices for accessing, scheduling, and paying for supercomputing power
- ✓ Hourly usage capability
  - alternative to weekly reservations
- ✓ Quicker access
  - advanced reservation not required to tap into capacity
- ✓ Flexible utilization management for client and IBM
  - Flexible capacity and resource sharing within and across sites

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Exa Corporation: DCCoD Client & Reseller

*Helping enable new capabilities in Automotive Computational Fluid Dynamics*

- **Scalable PowerFLOW® DWT™ (Digital Wind Tunnel)**
- **Exa® runs clients' aerodynamic, thermal, and aeroacoustic analysis using IBM DCCoD**
  - Ongoing IBM-Exa DCCoD partnership
  - Passenger cars, race cars, motorcycles, light trucks, heavy trucks
  - Over 15 Automotive clients served, over 20 current projects
- **Enabling new capabilities previously out of reach**
  - Large scale tests beyond the capacity of physical wind tunnels and in-house capacity
  - Easy capacity expansion for faster turnaround
  - More simulations and variants
  - More complex simulation models
  - More complex geometry and finer detail analysis
  - Larger data volumes
  - Support of projects that would otherwise forego analysis
  - More informed design decisions

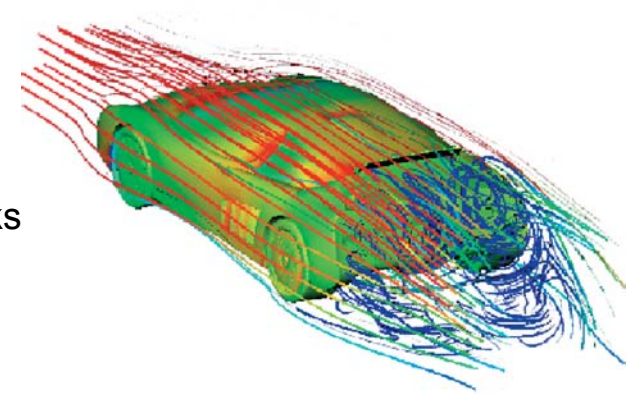

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Exa® Corporation: Customer Examples

- Scalability
  - ☑ From 4 CPUs in-house to 128 Opteron CPUs for 3 weeks
    - 32X capacity, more iterations, shorter design cycle
  - ☑ From 1 analysis iteration in 2 days to 10-50 runs overnight
- Quality & Time to Market
  - ☑ Support of projects that would otherwise forego analysis
  - ☑ More complex geometry and finer detail analysis
  - ☑ Multiple vehicle project reviews in one week, in-house hardware only capable of supporting one per week
- Multi-disciplinary optimization
  - ☑ Aerodynamic impact on suspension and driver fatigue
  - ☑ Aerodynamic impact on acoustics and comfort
  - ☑ Multi-vehicle studies for auto racing
- ☑ Earlier design cycle insights
  - ☑ Surface measurements (vel., temp., press.)
  - ☑ Underhood flow measurements (vel., temp
  - ☑ Underbody pressure losses
  - ☑ Environmental regulations
  - ☑ Acoustics noise sources and analysis
  - ☑ Brake cooling
  - ☑ More …

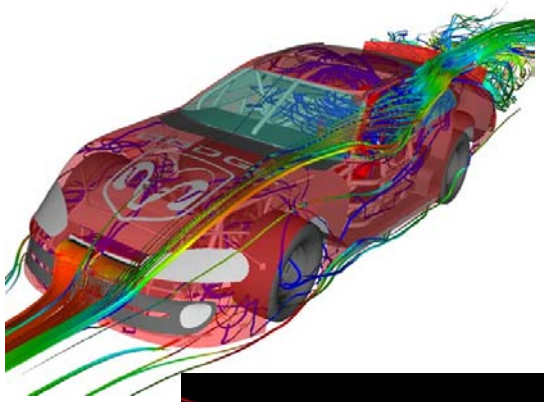

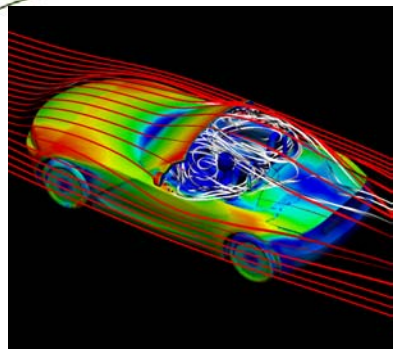

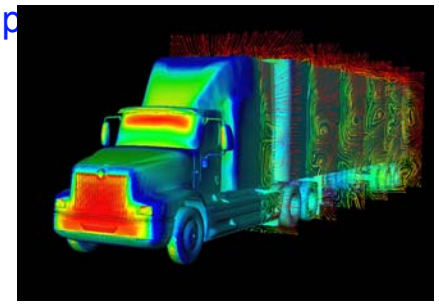

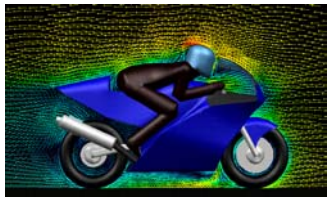

# IBM & Landmark Graphics Solutions

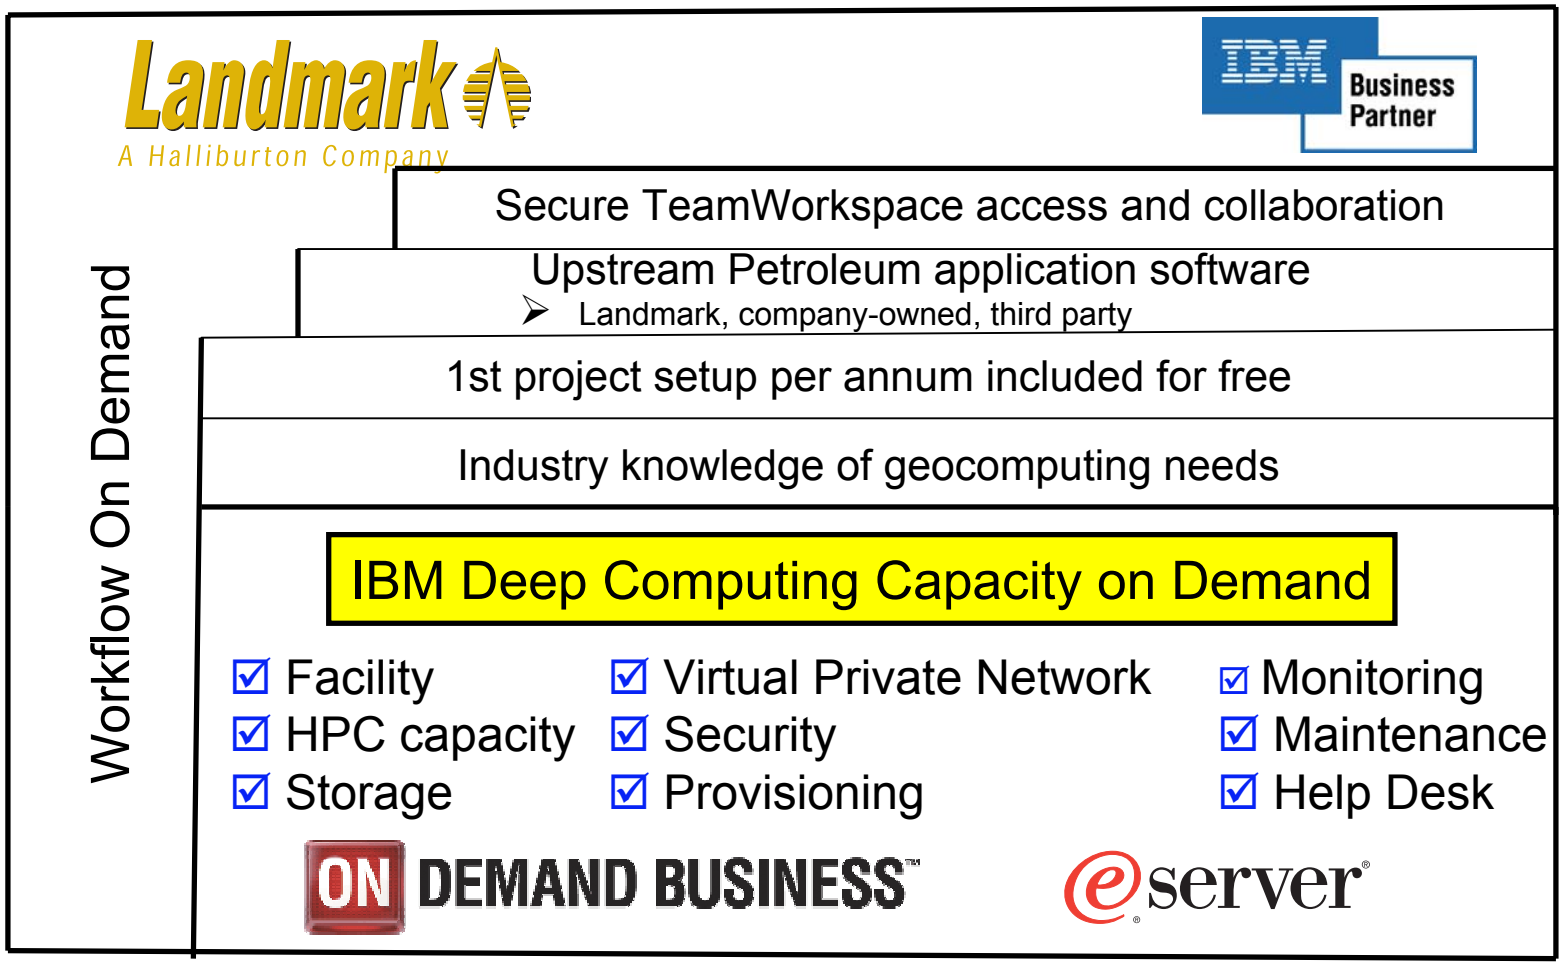


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# DCCoD Directions

- **Dynamic Scheduling & Provisioning Model (1Q05)**
- **Blue Gene, Cell, DCV**
- **Under evaluation**
  - Capacity growth in existing centers
  - New centers
  - Technology introductions
    - POWER5
    - BladeCenter JS20
    - Intel EM64T

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


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# The Future of High Performance Computing

Hybrid Computing

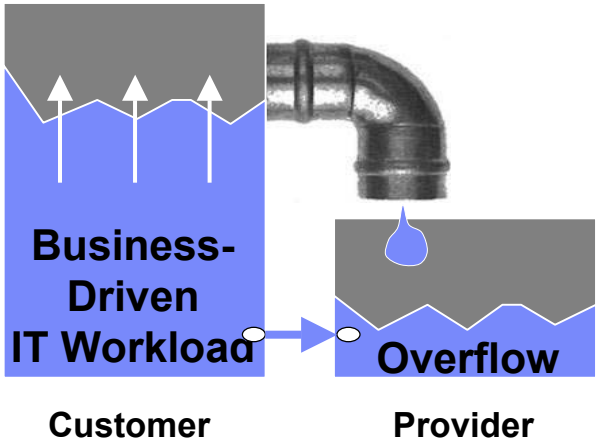


+

Grid Computing

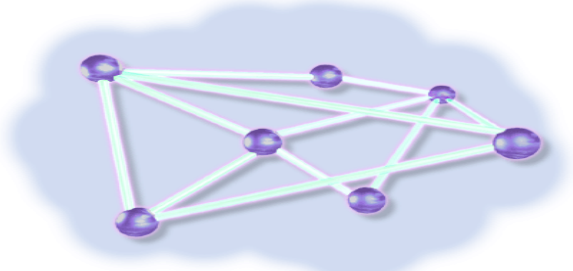

On Demand Computing

Utility Computing

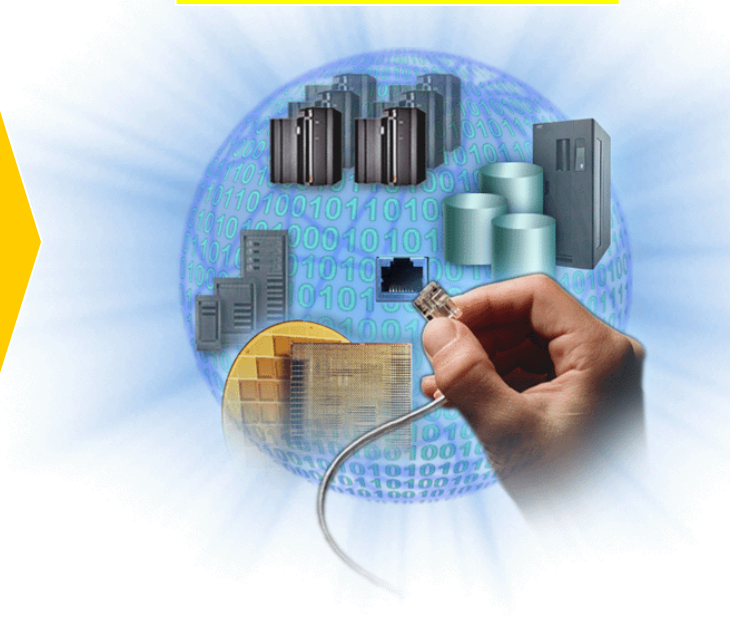

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

# Conclusion - Looking Forward

- **Operational trend to maximize available resources (people, capital, space, bandwidth, power, …)**
  - Fine-tuned capacity management (virtualization, partitioning)
  - Sophisticated networking and systems management
  - On-demand infrastructure – enabling businesses and research institutions to compete effectively on global level
- **Evolution in HPC software environment**
  - Emphasis on ease of use, simplified programming models
- **Intersection of "large and small"**
  - Information capture from multitude of pervasive/edge devices (e.g., RFID's, sensors) to central data stores and compute engines
  - Real-time mining and interpretation, integration across supply chain

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Selling Deep Computing Solutions – BG, Cell, DCV, DCCoD

- **Understand the opportunities**
  - Diverse set of markets with wide ranging needs for compute- and data-intensive workloads
  - Cuts across sectors, industries, geographies, company sizes, channels, …
  - These new Deep Computing offerings can set you apart from the competition
- **Know what to sell … seek help**
  - One size does not fit all – application "fit" and positioning is critical
    - "Mainstream" solution portfolio may be adequate
  - "Hybrid" transactions
    - Combine Linux cluster sale with DCCoD membership
    - Introduce DCV or DCCoD for differentiation and added value
  - Pre- and post-sales support - custom benchmarking, proofs of concept, centers of competency, ISV enablement, application porting and tuning, performance optimization, solutions assurance, competitive bids, skills transfer
  - IBM Research, E&TS
- **Contact your WW or geo Deep Computing sales experts**
  - Deep industry and technology expertise

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# SGB Seller Playbook

- **Available January 2005 from Systems Sales site**
- **Grid, Deep Computing, Linux and Virtualization Solutions**
- **Includes:**
  - Value propositions / key messages by industry
  - Qualification questions and next steps
  - Seller resources mapped to SSM steps
    - Customer success stories
    - Web casts, brochures, solutions briefs
    - Articles, press releases, analyst reports

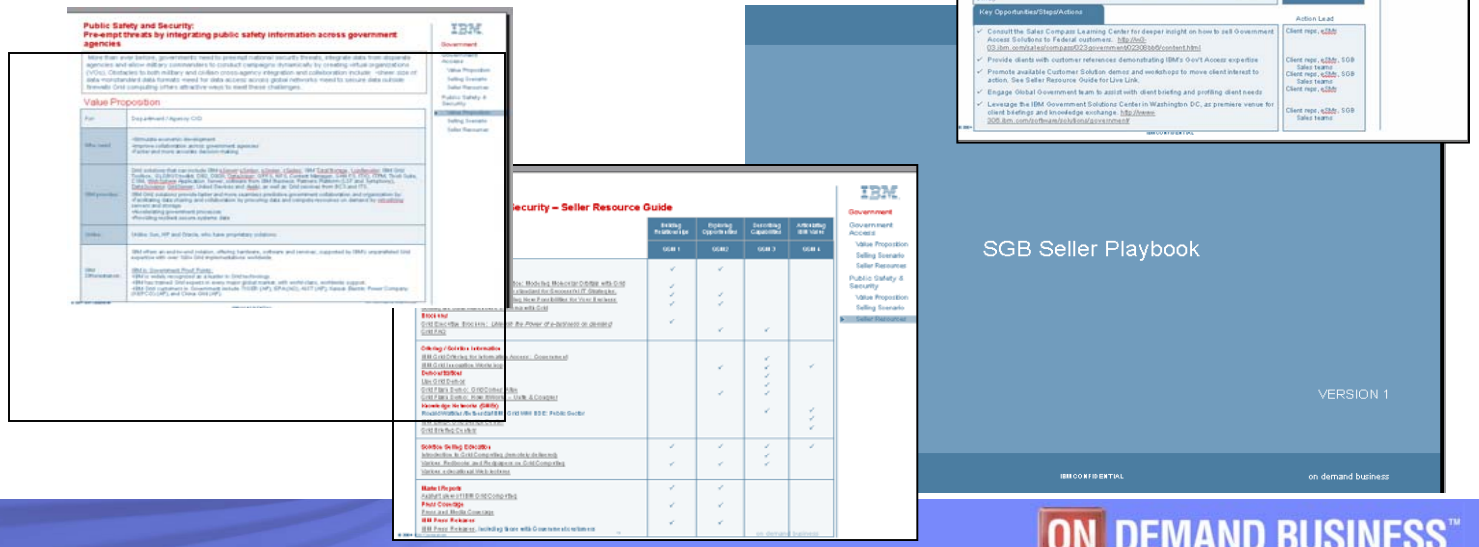


This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# Additional Resources

- http://www.ibm.com/servers/deepcomputing
- **Key contacts:**
  - WW Deep Computing Sales – Dave Jursik
  - WW DC Technical Sales Support – Kent Winchell
  - Americas – Joe Lopez
  - EMEA – Ian Green
  - AP – Sinisa Nikolic

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# End Of Presentation

This presentation is intended for the education of IBM and Business Partner sales personnel. It should not be distributed to customers.

# 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.

Many of the pSeries features described in this document are operating system dependent and may not be available on Linux. For more information, please check: http://www.ibm.com/servers/eserver/pseries/linux/whitepapers/linux_pseries.html

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.

Revised February 6, 2004

# 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, Blue Gene, Blue Lightning, C Set++, CICS, CICS/6000, CT/2, DataHub, DataJoiner, DB2, DEEP BLUE, developerWorks, DFDSM, DirectTalk, DYNIX, DYNIX/ptx, e business(logo), e(logo)business, e(logo)server, Enterprise Storage Server, ESCON, FlashCopy, GDDM, IBM, IBM(logo), ibm.com, IBM TotalStorage Proven, IntelliStation, IQ-Link, LANStreamer, LoadLeveler, Lotus, Lotus Notes, Lotusphere, Magstar, MediaStreamer, Micro Channel, MQSeries, Net.Data, Netfinity, NetView, Network Station, Notes, NUMA-Q, Operating System/2, Operating System/400, OS/2, OS/390, OS/400, Parallel Sysplex, PartnerLink, PartnerWorld, POWERparallel, PowerPC, PowerPC(logo), Predictive Failure Analysis, pSeries, PTX, ptx/ADMIN, RISC System/6000, RS/6000, S/390, Scalable POWERparallel Systems, SecureWay, Sequent, ServerProven, SP1, SP2, SpaceBall, System/390, The Engines of e-business, THINK, ThinkPad, 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/L(logo), AIX 5L, AIX PVMe, AS/400e, BladeCenter, Chipkill, Cloudscape, DB2 OLAP Server, DB2 Universal Database, DFDSM, DFSORT, Domino, e-business(logo), e-business on demand, eServer, GigaProcessor, HACMP, HACMP/6000, Hypervisor, i5/OS, IBMLink, IBM Virtualization Engine, IMS, Intelligent Miner, Micro-Partitioning, iSeries, NUMACenter, OpenPower, POWER, Power Architecture, Power Everywhere, PowerPC Architecture, PowerPC 603, PowerPC 603e, PowerPC 604, PowerPC 750, POWER2, POWER2 Architecture, POWER3, POWER4, POWER4+, POWER5, POWER5+, POWER6, Redbooks, Sequent (logo), SequentLINK, Server Advantage, ServeRAID, Service Director, SmoothStart, SP, S/390 Parallel Enterprise Server, ThinkVision, Tivoli Enterprise, TME 10, TotalStorage Proven, Ultramedia, VideoCharger, Visualization Data Explorer, X-Architecture, z/Architecture.

A full list of U.S. trademarks owned by IBM may be found at: http://www.**ibm.com**/legal/copytrade.shtml.

UNIX is a registered trademark in the United States, other countries or both.

Linux is a registered 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 and/or other countries.

Intel, Itanium and Pentium are registered trademarks and Intel Xeon and MMX are trademarks of Intel Corporation in the United States and/or other countries

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 and/or other countries.

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.

Other company, product and service names may be trademarks or service marks of others.

Revised August 23, 2004

# Notes on Benchmarks and Values

The benchmarks and values shown herein were derived using particular, well configured, development-level computer systems. Unless otherwise indicated for a system, the values were derived using external cache, if external cache is supported on the system. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application oriented testing. For additional information about the benchmarks, values and systems tested, contact your local IBM office or IBM authorized reseller or access the following on the Web:

| | | | |
|---|---|---|---|
| TPC | http://www.tpc.org | Linpack | http://www.netlib.no/netlib/benchmark/performance.ps |
| Pro/E | http://www.proe.com | SPEC | http://www.spec.org |
| GPC | http://www.spec.org/gpc | NotesBench Mail | http://www.notesbench.org |
| VolanoMark | http://www.volano.com | STREAM | http://www.cs.virginia.edu/stream/ |

Unless otherwise indicated for a system, the performance benchmarks were conducted using AIX V4.3 or AIX 5L. IBM C Set++ for AIX and IBM XL FORTRAN for AIX with optimization were the compilers used in the benchmark tests. The preprocessors used in some benchmark tests include KAP 3.2 for FORTRAN and KAP/C 1.4.2 from Kuck & Associates and VAST-2 v4.01X8 from Pacific-Sierra Research. The preprocessors were purchased separately from these vendors. Other software packages like IBM ESSL for AIX and MASS for AIX were also used in some benchmarks.

The following SPEC and Linpack benchmarks reflect microprocessor, memory architecture, and compiler performance of the tested system (XX is either 95 or 2000):

- SPECintXX - SPEC component-level benchmark that measures integer performance. Result is the geometric mean of eight tests comprising the CINTXX benchmark suite. All of these are written in the C language. SPECint_baseXX is the result of the same tests as CINTXX with a maximum of four compiler flags that must be used in all eight tests.
- SPECint_rateXX - Geometric average of the eight SPEC rates from the SPEC integer tests (CINTXX). SPECint_base_rateXX is the result of the same tests as CINTXX with a maximum of four compiler flags that must be used in all eight tests.
- SPECfpXX - SPEC component-level benchmark that measures floating-point performance. Result is the geometric mean of ten tests, all written in FORTRAN, included in the CFPXX benchmark suite. SPECfp_baseXX is the result of the same tests as CFPXX with a maximum of four compiler flags that must be used in all ten tests.
- SPECfp_rateXX - Geometric average of the ten SPEC rates from SPEC floating-point tests (CFPXX). SPECfp_base_rateXX is the result of the same tests as CFPXX with a maximum of four compiler flags that must be used in all ten tests.
- SPECweb96 - Maximum number of Hypertext Transfer Protocol (HTTP) operations per second achieved on the SPECweb96 benchmark without significant degradation of response time. The Web server software is ZEUS v.1.1 from Zeus Technology Ltd.
- SPECweb99 - Number of conforming, simultaneous connections the Web server can support using a predefined workload. The SPECweb99 test harness emulates clients sending the HTTP requests in the workload over slow Internet connections to the Web server. The Web server software is Zeus from Zeus Technology Ltd.
- SPECweb99_SSL - Number of conforming, simultaneous SSL encryption/decryption connections the Web server can support using a predefined workload. The Web server software is Zeus from Zeus Technology Ltd.
- SPEC OMP2001 - Measures performance based on OpenMP applications.
- SPECsfs97_R1 - Measures speed and request-handling capabilities of NFS (network file server) computers.

Revised August 23, 2004

# Notes on Benchmarks and Values (Cont.)

- SPECjAppServer200X (where X is 1 or 2) - Measures the performance of Java Enterprise Application Servers using a subset of J2EE APIs in a complete end-to-end Web application.

The Linpack benchmark measures floating-point performance of a system.

- Linpack DP (Double Precision) - n=100 is the array size. The results are measured in megaflops (MFLOPS).
- Linpack SP (Single Precision) - n=100 is the array size. The results are measured in MFLOPS.
- Linpack TPP (Toward Peak Performance) - n=1,000 is the array size. The results are measured in MFLOPS.
- Linpack HPC (Highly Parallel Computing) - solves the largest system of linear equations possible. The results are measured in GFLOPS.

STREAM is a simple synthetic benchmark program that measures sustainable memory bandwidth (in MB/s) and the corresponding computation rate for simple vector kernels. Both standard and tuned results may be reported. http://www.cc.virginia.edu/stream/

VolanoMark is a 100% pure Java server benchmark that creates long-lasting network client connections in groups of 20 and measures how long it takes for the clients to take turns broadcasting their messages to the group. The benchmark reports a score as the average number of messages transferred by the server per second.

- The following Transaction Processing Performance Council (TPC) benchmarks reflect the performance of the microprocessor, memory subsystem, disk subsystem, and some portions of the network:
  - tpmC - TPC Benchmark C throughput measured as the average number of transactions processed per minute during a valid TPC-C configuration run of at least twenty minutes.
  - $/tpmC - TPC Benchmark C price/performance ratio reflects the estimated five year total cost of ownership for system hardware, software, and maintenance and is determined by dividing such estimated total cost by the tpmC for the system.
  - QppH is the power metric of TPC-H and is based on a geometric mean of the 17 TPC-H queries, the insert test, and the delete test. It measures the ability of the system to give a single user the best possible response time by harnessing all available resources. QppH is scaled based on database size from 30GB to 10TB.
  - QthH is the throughput metric of TPC-H and is a classical throughput measurement characterizing the ability of the system to support a multiuser workload in a balanced way. A number of query users is chosen, each of which must execute the full set of 17 queries in a different order. In the background, there is an update stream running a series of insert/delete operations. QthH is scaled based on the database size from 30GB to 10TB.
  - $/QphH is the price/performance metric for the TPC-H benchmark where QphH is the geometric mean of QppH and QthH. The price is the five-year cost of ownership for the tested configuration and includes maintenance and software support.

Revised August 23, 2004

# Notes on Benchmarks and Values (Cont.)

The following graphics benchmarks reflect the performance of the microprocessor, memory subsystem, and graphics adapter:

- SPECxpc results - Xmark93 is the weighted geometric mean of 447 tests executed in the x11perf suite and is an indicator of 2D graphics performance in an X environment. Larger values indicate better performance.
- SPECplb results (graPHIGS) - PLBwire93 and PLBsurf93 are geometric means of literal and optimized Picture Level Benchmark (PLB) tests for 3D wireframe and 3D surface tests, respectively. Larger values indicate better performance.
- SPECopc results - Viewperf 7 (3dsmax-01, DRV-08, DX-07, Light-05, ProE-01, UGS-01) and Viewperf 6.1.2 (AWadvs-04, DRV-07, DX-06, Light-04, medMCAD-01, ProCDRS-03) are weighted geometric means of individual viewset metrics. Larger values indicate better performance.

The following graphics benchmarks reflect the performance of the microprocessor, memory subsystem, graphics adapter and disk subsystem.

- SPECapc Pro/Engineer 2000i2 results - PROE2000I2_2000370 was developed by the SPECapc committee to measure UNIX and Windows workstations in a comparable real-world environment. Larger numbers indicate better performance.

The NotesBench Mail workload simulates users reading and sending mail. A simulated user will execute a prescribed set of functions 4 times per hour and will generate mail traffic about every 90 minutes. Performance metrics are:

- NotesMark - transactions/minute (TPM).
- NotesBench users - number of client (user) sessions being simulated by the NotesBench workload.
- $/NotesMark - ratio of total system cost divided by the NotesMark (TPM) achieved on the Mail workload.
- $/User - ratio of total system cost divided by the number of client sessions successfully simulated for the NotesBench Mail workload measured. Total system cost is the price of the server under test to the client, including hardware, operating system, and Domino Server licenses.

Application Benchmarks

- SAP - Benchmark overview information: http://www.sap-ag.de/solutions/technology/bench.htm; Benchmark White Paper September, 2000; http://www.sap-ag.de/solutions/technology/pdf/50020428.pdf.
- PeopleSoft - To get information on PeopleSoft benchmarks, contact PeopleSoft directly or the PeopleSoft/IBM International Competency Center in San Mateo, CA.
- Oracle Applications - Benchmark overview information: http://www.oracle.com/apps_benchmark/results/results.html
- Baan - The Baan benchmark demonstrates the scalability of Baan ERP solutions. The test results provide the number of Baan Reference Users (BRUs) that can be supported on a specific system. BRU is a single on-line user or a batch unit workload. These metrics are consistent with those used internally by both IBM and Baan to size systems. To get more information on Baan benchmarks, go to http://www.ssaglobal.com.
- J.D. Edwards Applications - Product overview information at http://www.jdedwards.com.

Revised August 23, 2004

# Notes on Performance Estimates

rPerf

- rPerf (Relative Performance) is an estimate of commercial processing performance relative to other pSeries systems. It is derived from an IBM analytical model which uses characteristics from IBM internal workloads, TPC and SPEC benchmarks. The rPerf model is not intended to represent any specific public benchmark results and should not be reasonably used in that way. The model simulates some of the system operations such as CPU, cache and memory. However, the model does not simulate disk or network I/O operations.
- rPerf estimates are calculated based on systems with the latest levels of AIX 5L and other pertinent software at the time of system announcement. Actual performance will vary based on application and configuration specifics. The IBM @server pSeries 640 is the baseline reference system and has a value of 1.0. Although rPerf may be used to approximate relative IBM UNIX commercial processing performance, actual system performance may vary and is dependent upon many factors including system hardware configuration and software design and configuration.
- All performance estimates are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, and application sizing guides to evaluate the performance of a system they are considering buying. For additional information about rPerf, contact your local IBM office or IBM authorized reseller.

Revised June 28, 2004