IBM Systems and Technology Group University 2005

San Diego, California
January 16 - 19, 2005

ON DEMAND BUSINESS™

v 1.0

IBM Systems and Technology Group University 2005

Deep Computing Overview: IBM eServer and TotalStorage

Course #: CB51

Rebecca Austen
Director, Deep Computing Marketing

v 1.0

Agenda

- **Deep Computing Overview**
  - Application, Technology and Market Trends
  - Deep Computing Strategy
  - Portfolio and Positioning
- **Deep Computing Industry Solutions**
  - Industrial/Automotive
  - Digital Media
  - Life Sciences
  - Business Intelligence
- **Conclusion**
- **Resources**

Learning Objectives

- Understand the Deep Computing market place
- Identify Deep Computing opportunities in your industry
- Leverage the STG product line and Deep Computing solutions portfolio to win business
- Know whom to contact and where to find more information on Deep Computing

Deep Computing Overview

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

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

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

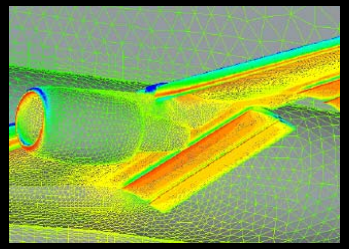

CFD Wing Simulation

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

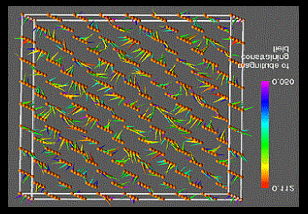

Magnetic Materials:

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

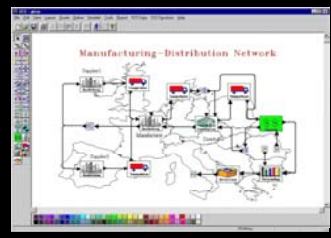

Source: B. Dietrich, IBM

Deep Computing: Energy Production and Trading

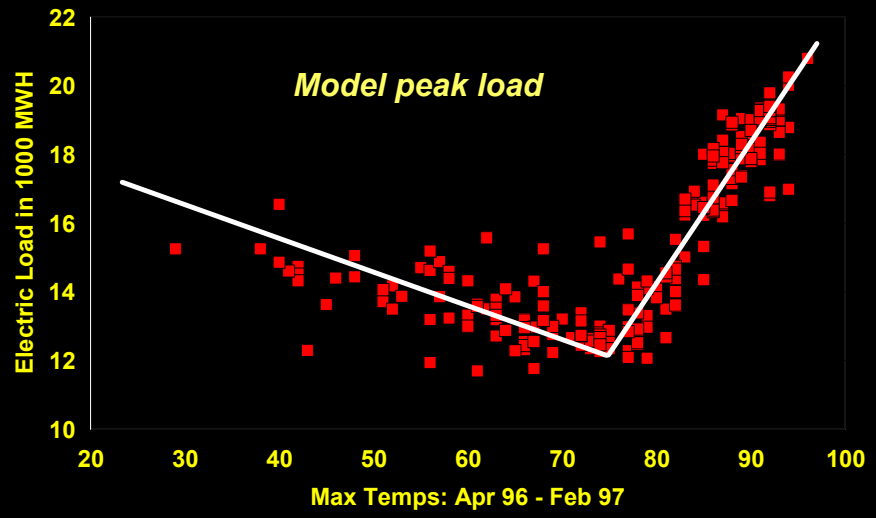

The application of computational methods
- •Mathematical modeling
- •Optimization
- •Statistics
- •Design and analysis of algorithms
- •Parallel computing

to large data sets
- •Historical
- •Operational

to solve business decision problems

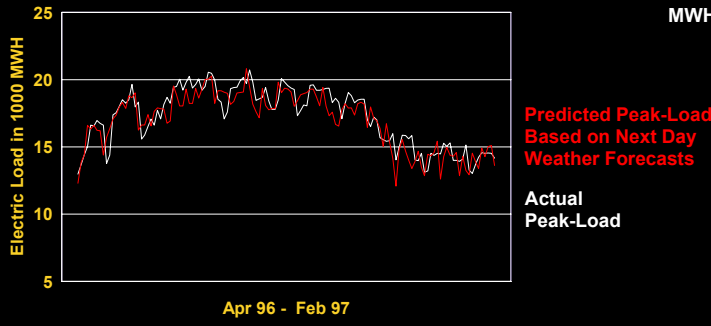

Predict peak load based on weather forecast

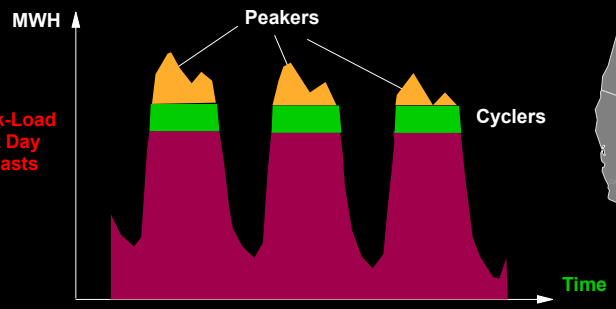

Plan production schedules

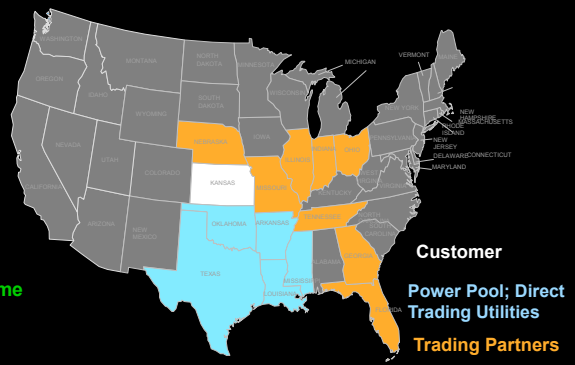

Identify energy trading partners

Source: W. Pulleyblank, IBM

Market Opportunity & Growth

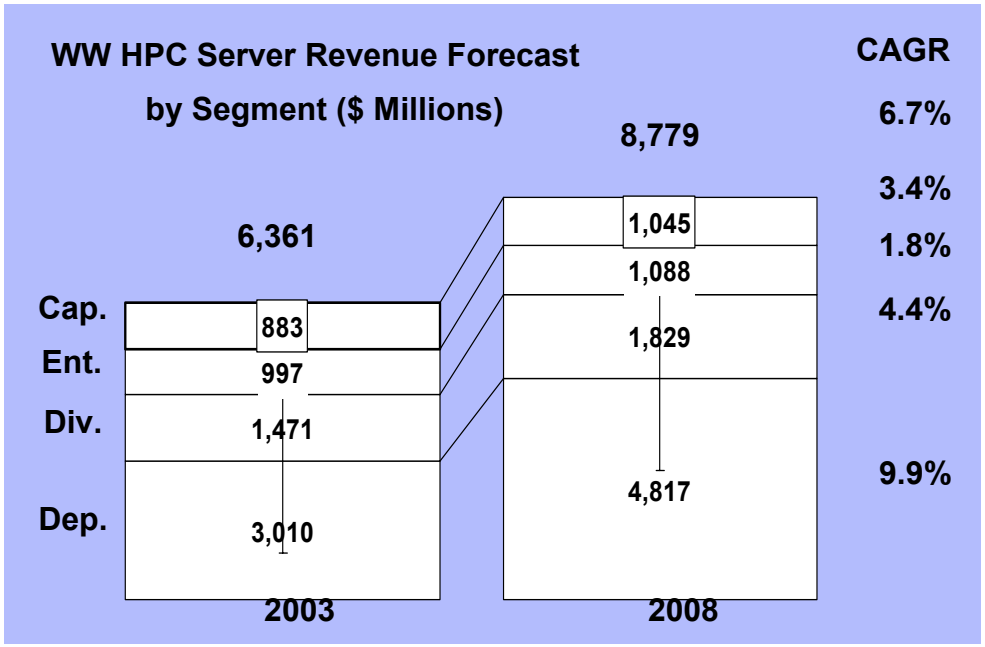

Source: IDC Cluster Multi-Client Study, May 2004

- 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

HPC Server Revenue and Share

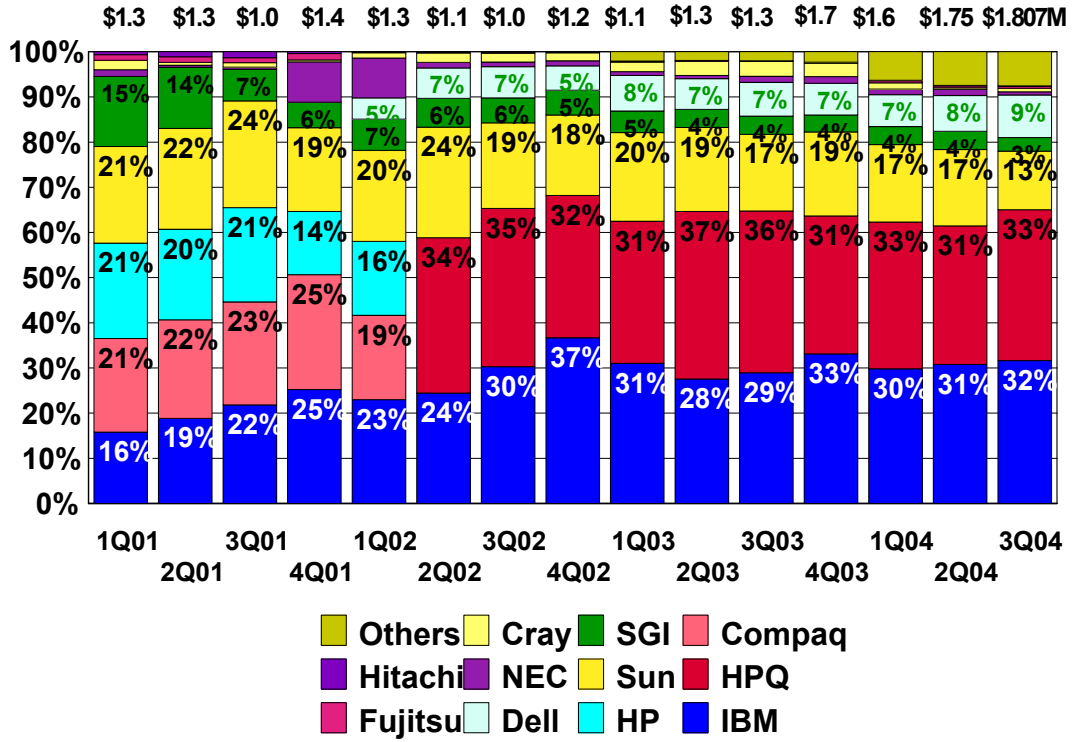

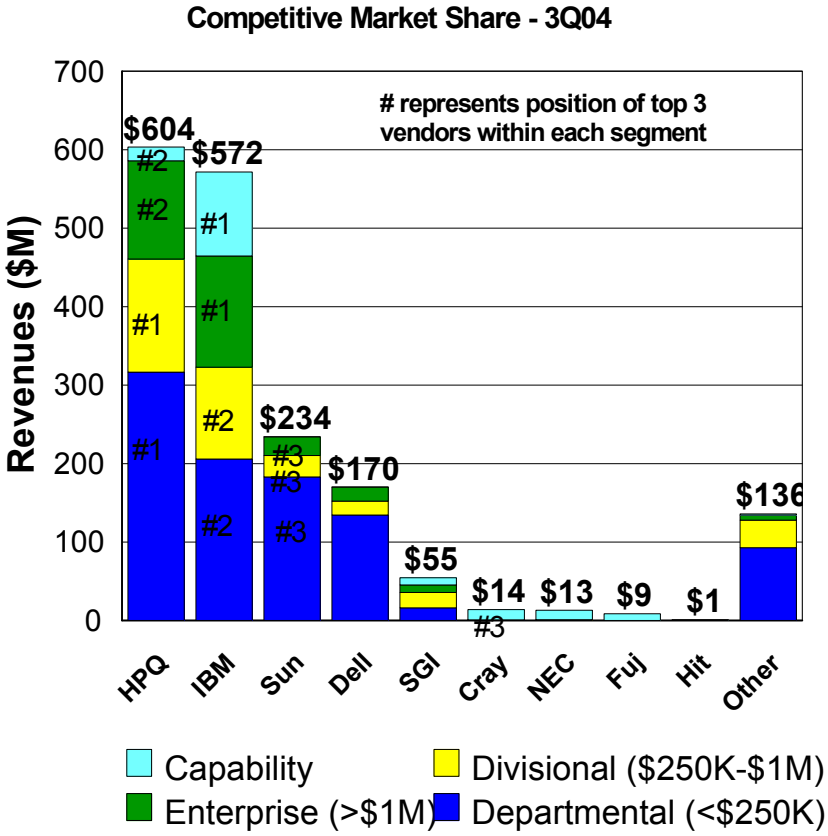

Source: IDC 12/2004

HPC Operating System Trends

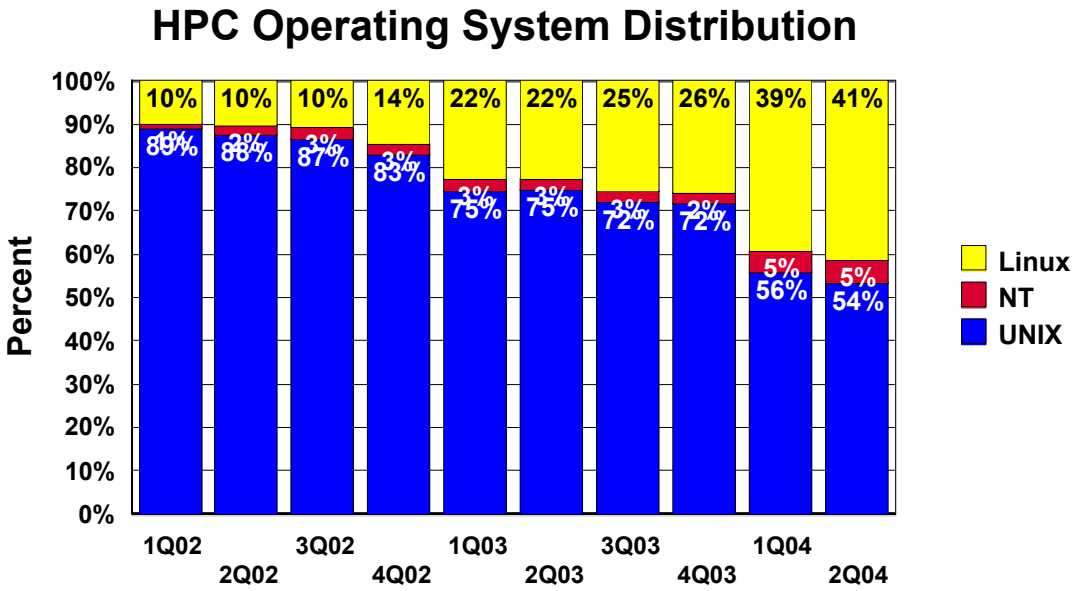

Distribution of Operating System by IDC Segment - 2Q04

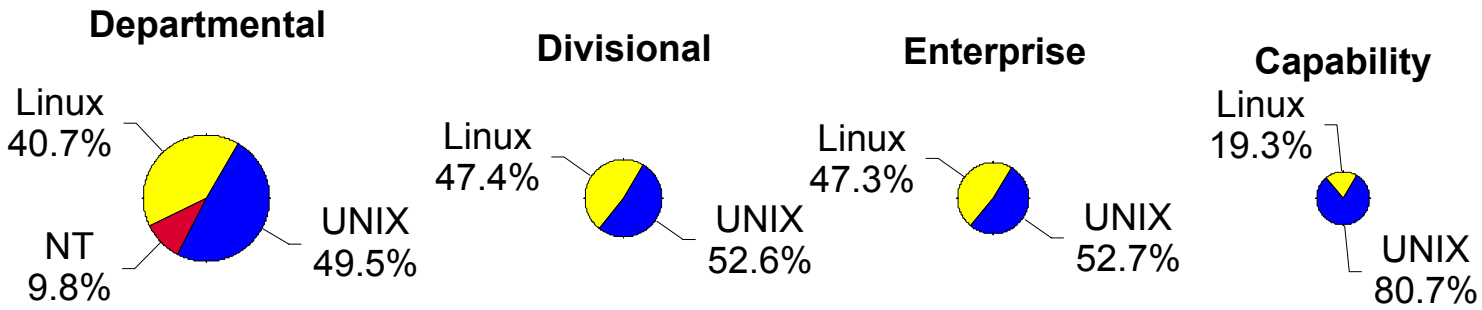

Source: IDC 09/2004

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%)

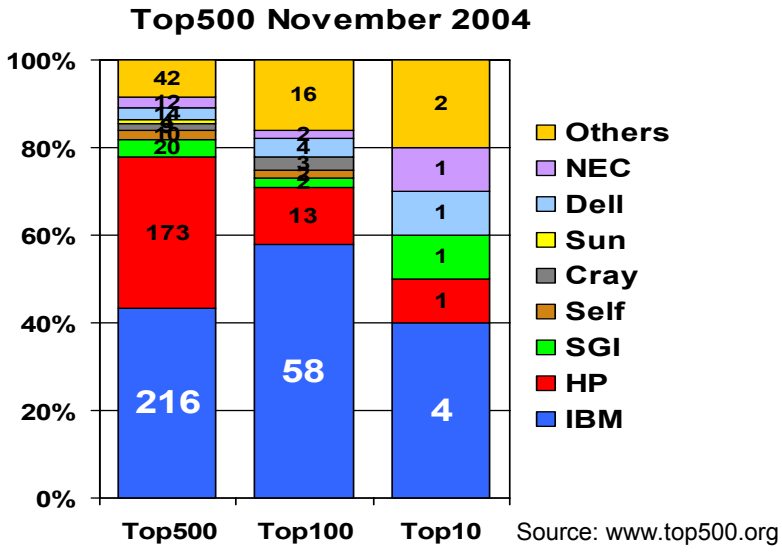

Source: www.top500.org

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

2U/8-16w

Switch

SCALE IN

1U/1-2w

Switch

SCALE OUT

"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**

Deep Computing

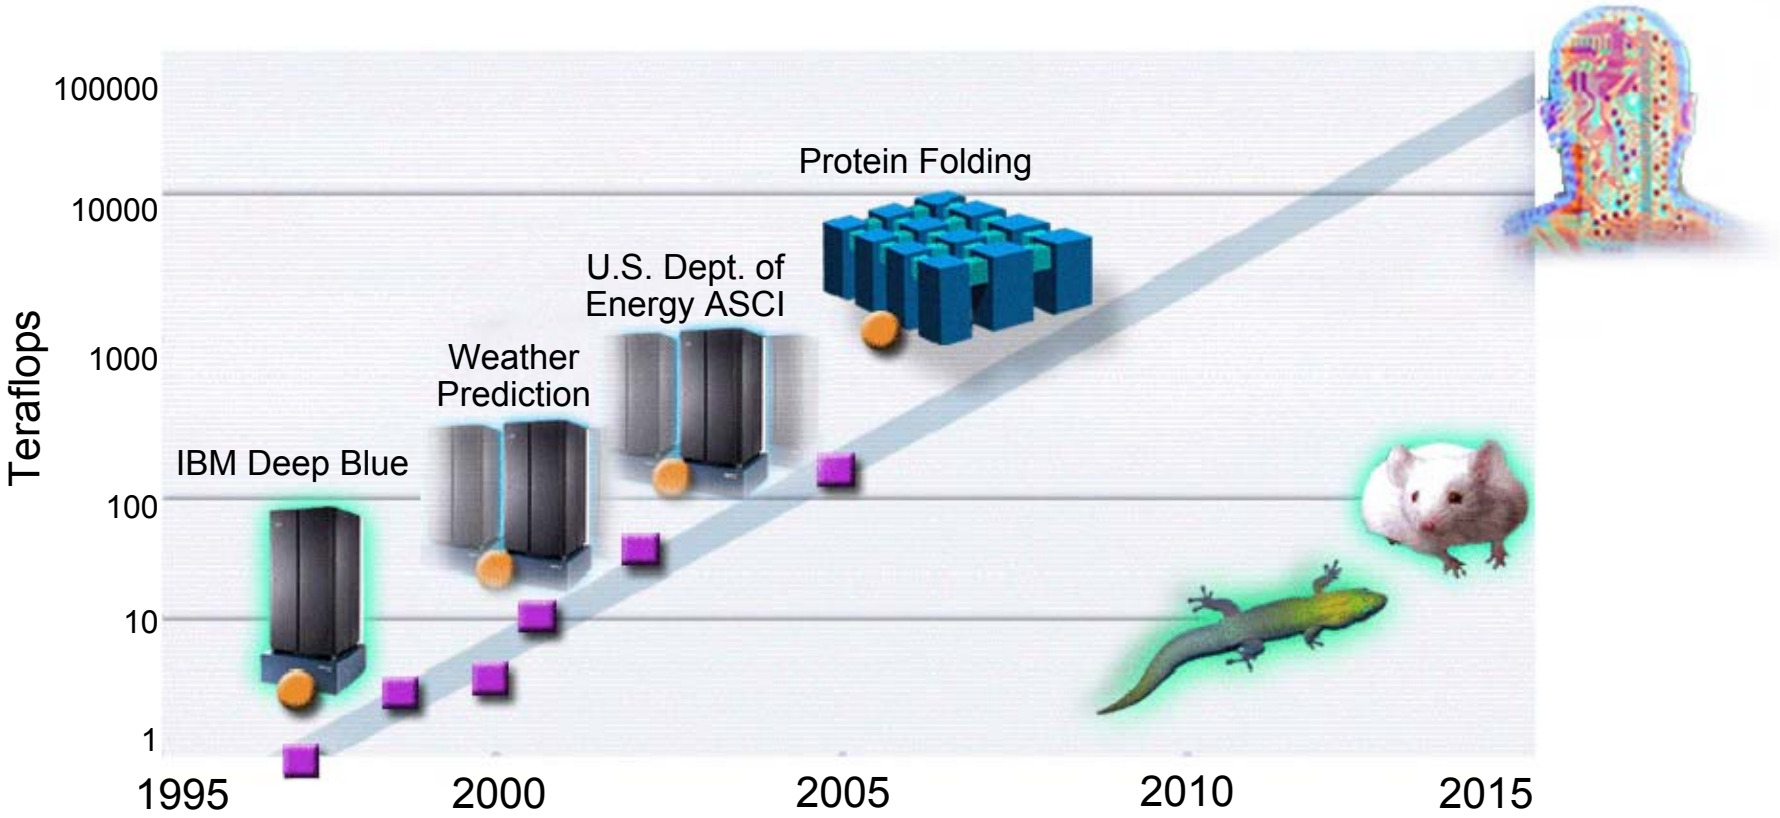

Deep Computing Leadership Initiatives

- **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**

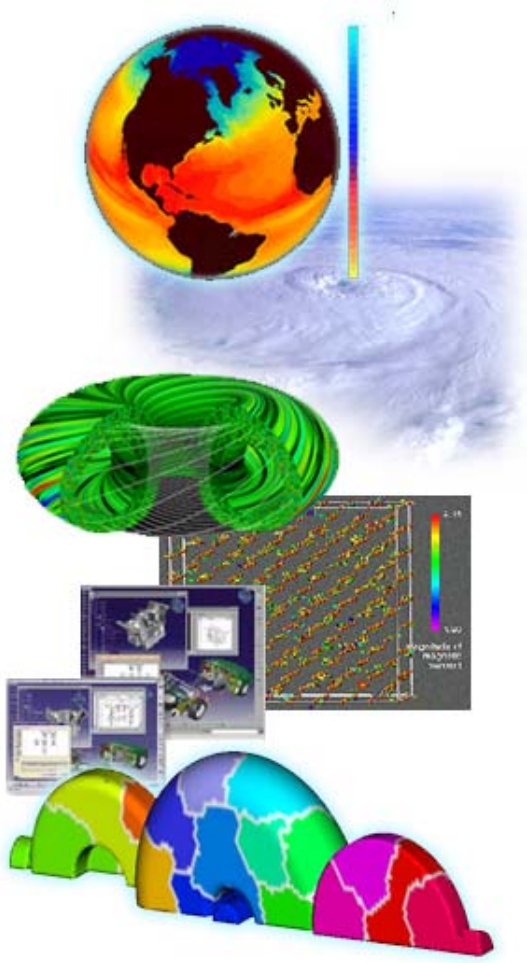

IBM's Deep Computing Strategy

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**

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

IBM Systems – Industry Leadership and Choice

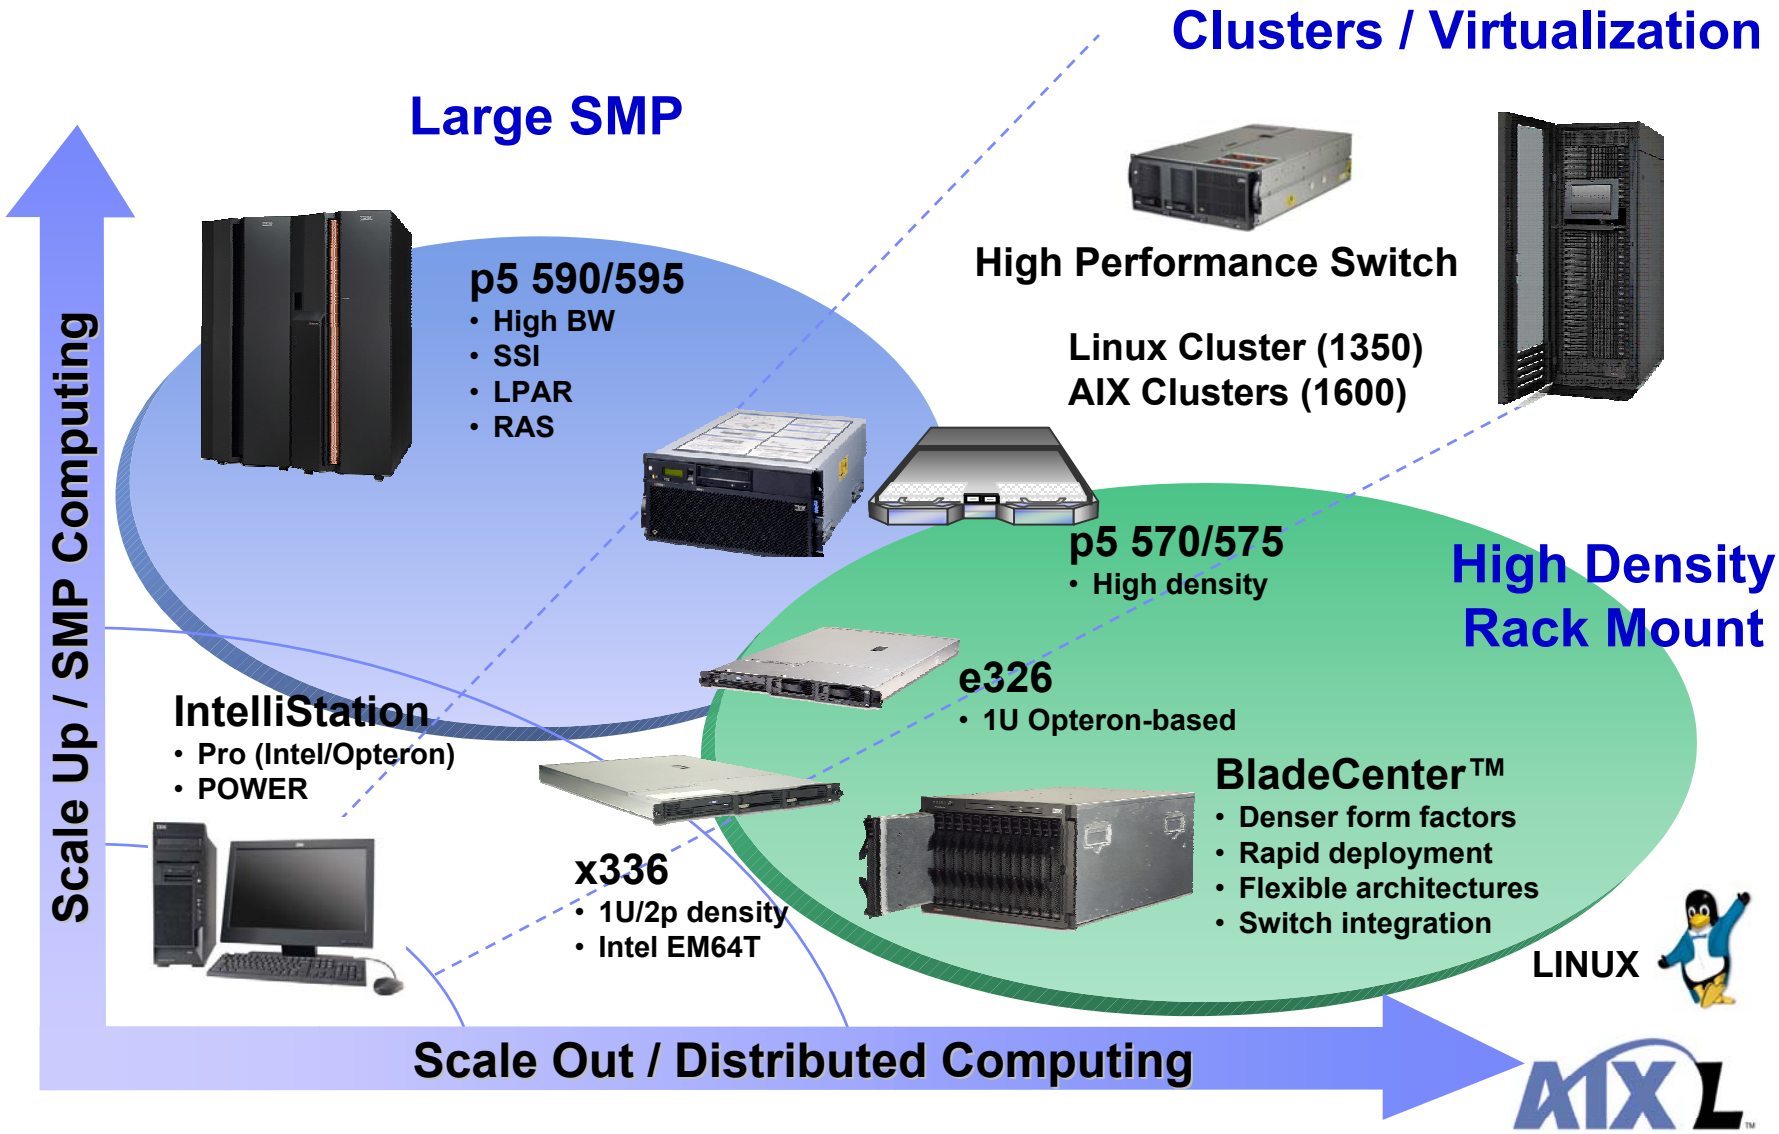

HPC Server Market by Deployment Model

Scale-out (high-value)	Purpose-built
Tightly coupled clusters of RISC or Itanium 2-based SMP servers; industry standard or custom packaging; vendor integrated; (optional) high performance interconnect (industry standard, OEM or custom) IBM Offerings: p655, p690, p5 575	Usually custom microprocessors, usually employ vectors and streaming, custom interconnect, custom packaging, vendor integration, specifically designed for HPC capability workloads IBM Offerings: Off-roadmap offerings like Blue Gene via special bid
Scale-out (commodity)	**(Standalone) SMP**
Clusters of 1, 2, and 4-way blade or rack-optimized servers, based on “merchant” or low-cost technology, standard or OEM interconnects, standard packaging, often vendor integrated IBM Offerings: BladeCenter HS20 / JS20; e325/326, x335/336, x345/346, x382	2-way to 64-way (or bigger) microprocessor-based SMP servers IBM Offerings: e325/326, x335/336, p615, p630, p650, p670, p690, p5-520, 550, 570, 590, 595

Source: IDC 1Q04 and team analysis

Workload Positioning … One Size Doesn't Fit All

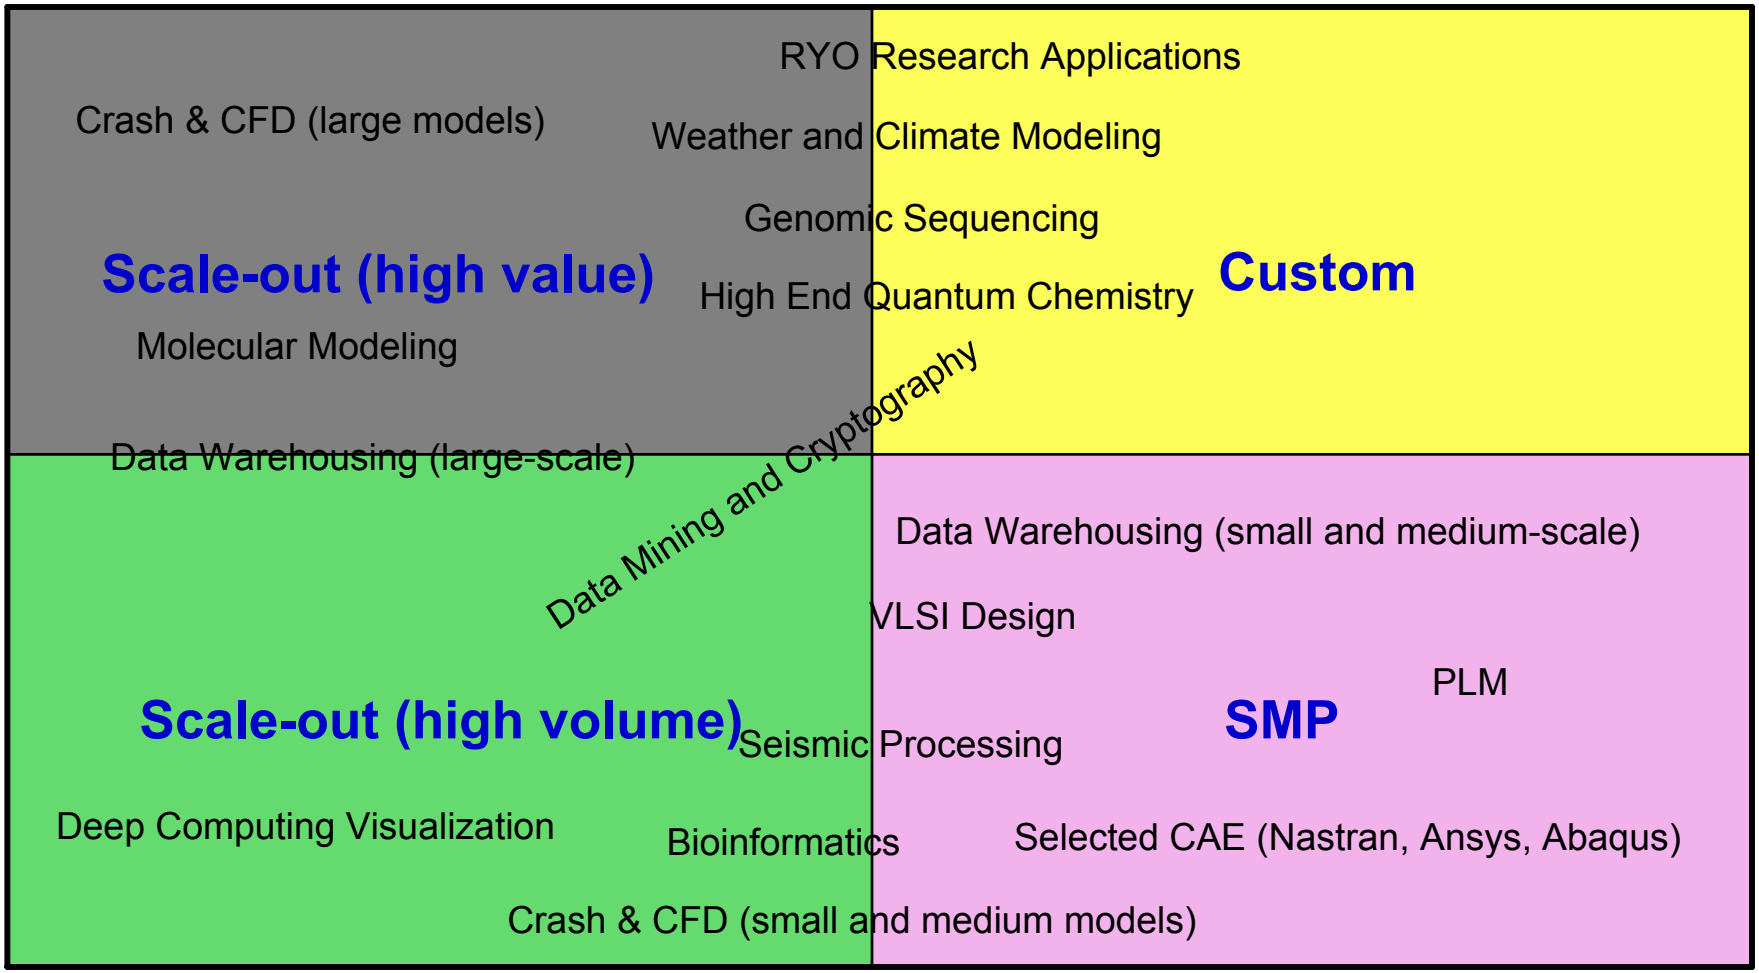

64-BIT – An Evolving Marketplace Architecture

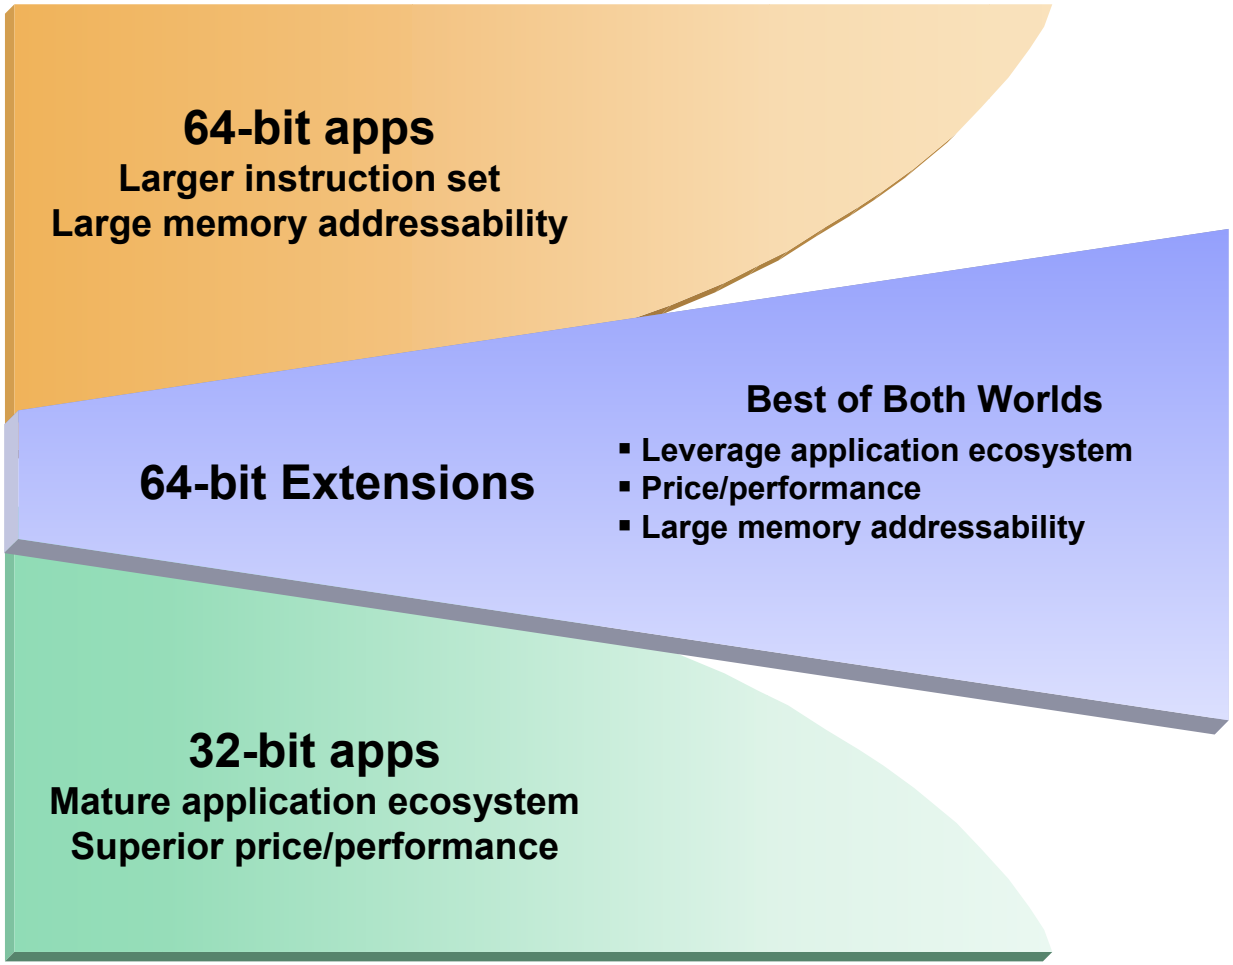

AMD Opteron and Intel EM64T

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

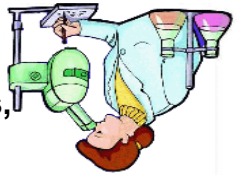

Business Intelligence

Data warehousing and data mining

Financial Services

Optimizing IT infrastructure, risk management and compliance, analytics

Petroleum

Oil and gas exploration and production

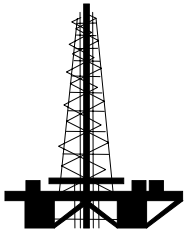

Industrial/Product Lifecycle Management

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

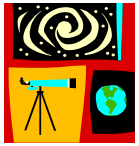

Government & Higher Education

Scientific research, classified/defense, weather/environmental sciences

Deep Computing Application Portfolio

Life Science

Discovery
- ★Accelrys*
- ★CCG*
- ★Spotfire
- + Oracle- Cross
- + Biomax- Geo
- + Platform Comp*-Grid
- + ThermoElectron(L/P)
- + Tripos
- + Wavefunction

Development
- ★Waters Corp (NuGen+Micromass)
- ★SAS- Cross
- ★Siebel - Cross
- ★Documentum- Cross
- + PhaseForward

Med Imaging/Clinical
- ★DeCode
- ★Cerner* (L/P)
- + Emageon
- + Seimens
- + Mckesson- SMB AG
- + Philips
- + Agfa
- + Kodak
- + GE
- + Phillips
- + Fuji
- + Stentor
- ~~Gaussian*~~
- ✓Daylight (recruit)
- ✓MDL* (recruit)

- ✓Apocon
- ✓Applied BioSys*
- ✓CAChe*
- ✓CCP4
- ✓Brunel Univ*
- ✓Cambridge Chrystallo*
- ✓CambridgeSoft
- ✓Compugen
- ✓CSD
- ✓EBI*
- ✓EMBL
- ✓Electri Genetics
- ✓GA Tech*
- ✓Gene IT
- ✓GeneData AG*
- ✓Geospiza
- ✓Global Phasing Ltd
- ✓Harvard/Scripps
- ✓HGMP-RC
- ✓Hutchinson Cancer Resh
- ✓IDBS*
- ✓Iowa State U*
- ✓Los Alamos Nat'l Lab*
- ✓Matrix Sciences*
- ✓MIT
- ✓NCBI*
- ✓NCGR*
- ✓NuTec Sciences*
- ✓Open Eye Sci. SW
- ✓Optive Research
- ✓PNNL*
- ✓Proteome
- ✓QChem*
- ✓Rockefeller University

- ✓Royal Institute*
- ✓Roseta
- ✓Sanger*
- ✓Schrodinger*
- ✓Scripps Research.
- ✓SDSC*
- ✓Secant
- ✓Semichem*
- ✓Silicon Genetics
- ✓Southwest Parallel
- ✓SPSoft*
- ✓Stranford*
- ✓Syrxx*
- ✓Thistlesoft*
- ✓Turboworx
- ✓UC at Berkley*
- ✓UCSF *
- ✓USC *
- ✓Univ of Leeds *
- ✓Univ of Minnesota *
- ✓Univ of Ottawa *
- ✓Univ of Pitt *
- ✓Univ of Vienna *
- ✓Univ of Virginia *
- ✓Univ of Birmingham*
- ✓Univ of Washington*
- ✓Vrije Univ. *
- ✓Wash U (St Louis) *
- ✓Whitehead Institute*

Digital Media

Digital Content Creation (DCC)
- • Alias
- • Discreet
- • Avid
- • Softimage

Online Games
- +Criterion
- +Havoc

Digital Media Center (DMC)
- +FP Digital
- +Masstech
- +Apple
- +Pinnacle
- +Thomson GV

Radio Asset Mgt (RAM)
- +Dalet
- +Jutel

Broadcast Asset Mgt (BAM)
- + Ardendo
- + Napster (AG)
- + David (EMEA)
- + Dayang (AP)

Digital Content Mgt (DCM)
- + Ancept

Digital Security /Surveilance
- ★ESRI
- + Verint
- + Nice

Video on Demand (VOD)
- + Orca
- + Entone
- + Kasena

Petroleum

- + Landmark Graph
- + Schlumberger/SIS
- + Paradigm Geophy.
- + CGG (L/P)
- + PGS (L/P)
- ✓CMG Ltd
- ✓Beicip-Franlab
- ✓GX Technology
- ✓WesternGeco
- ✓Earth Decisions Sciences
- ° Data Modeling
- ° GeoCenter
- ° Veritas DGC
- ° ERSI
- ° Telegnomics

PLM

CAD/PDM
- ★Dassault Systemes
- + UGS
- + PTC
- + MatrixOne
- ★SAP PLM
- • Autodesk
- • SolidWorks

EDA
- +Cadence
- +Mentor Graphics
- +Synopsys
- ~~Qad- SMB AG~~
- ° Zuken (Geo)
- ° Magma Design Automation
- ° Metacomp
- ° Silvaco Int'l

CAE
- ★MSC*
- + ABAQUS *
- + Fluent*
- + ESI
- + Ansys/(CEI/ICEM)*
- + ESI
- + CD adapco Grp *
- + CEI (Viizualiz.)*
- + LSTC*
- + Mecalog *
- ★Engenious
- ✓Altair
- ✓AVL *
- ✓CDH *
- ✓EXA *
- ~~✓LMS *~~
- ✓NASA Aero*
- ✓TNO Madymo
- ✓Ricardo SW
- ✓UT (AMLS)*
- ° Adina, Inc.
- ° Algor
- ° Analyt. Meth*
- ° Autoform Eng.
- ° CADSI
- ° Boeing BCLIB-EXT
- ° ETA
- ° iLOG
- ° LANL (Kiva)*
- ° Moldflow *
- ° Multiport
- ° SRAC *

HPC FSS/BI

- ★SearchSpace
- ★Algorithmics
- ★SAS
- ✓Axion
- ✓Fair Issac

Public Domain/ RYO

- 4DVAR *
- Arpege/ Aladin *
- ARPS
- dGauss *
- ECHAM*
- Eta *
- fetchdom
- HIRLAM *
- IFS *
- LM *
- LM *
- MOM *
- MM5 *
- POP *
- SQ *
- SVD *
- Topogen *
- Unifed - Model*
- ★SVD *

Tools/Research

- ✓Avaki *
- ✓Applied Parallel Research *
- ✓Data Synapse
- ✓Eng. Intelli
- ✓Etnus *
- ✓GNU *
- ✓IBM's Rational*
- ✓Intel's KAI *
- ✓NAG, Ltd *
- ✓PathScale*
- ✓Scali*
- °Argone Lab*
- °CEBPA
- °Geodesic *
- °MatrixX *
- °MP/SW Tech *
- °Pacific Sierra *
- °Pallas *
- °Portland Grp *
- °Rice University*
- °Tallas
- °Wayne State Univ. *
- °Univ. of Arkansas *

Legend

- ★Strategic Alliance- ISG+Sector+DR Cov.
- ★ISV Advantage- ISG+Sector Coverage
- + ISG ISV Coverage
- ✓ISG Technical Coverage
- • Covered by x/IntelliStation Org
- ° Network Coverage required
- *Public Sector ISVs/Codes

High Priority

Deep Computing Industry Solution Examples

- **Industrial/Automotive**
- **Digital Media**
- **Life Sciences**
- **Business Intelligence**

IBM Automotive Engineering Innovation Framework
Delivering an integrated innovation infrastructure

IBM is building a set of world class offerings comprised of the best of IBM technologies, systems, storage, visualization offerings and infrastructure, built on a common architectural framework, integrated and optimized in partnership with leading Automotive ISVs

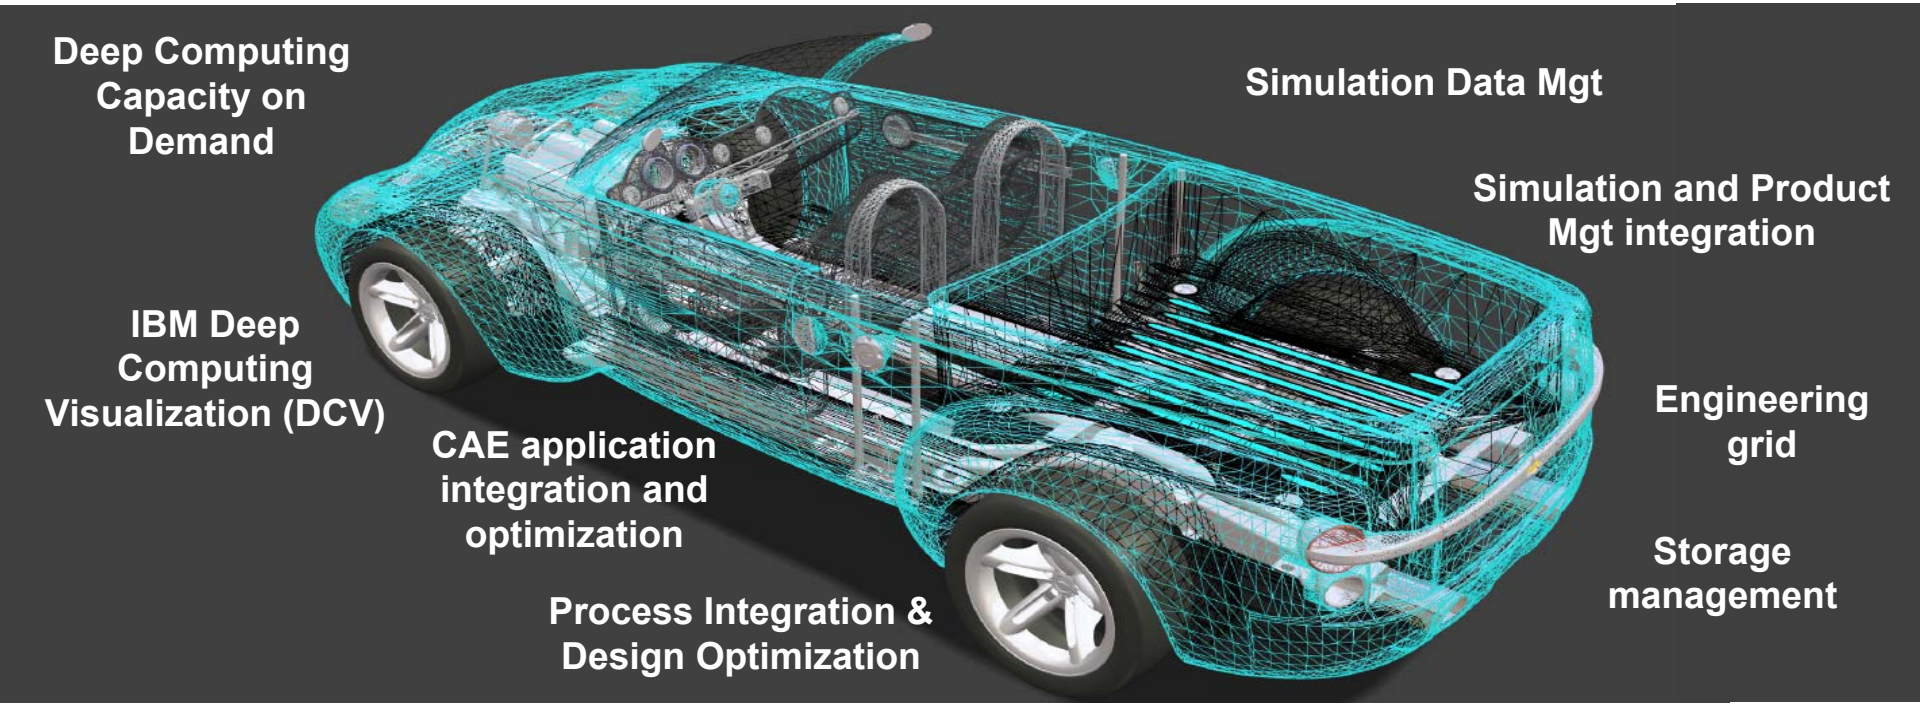

Analysis and simulation (CAE) is essential component of the Integrated Digital Prototype

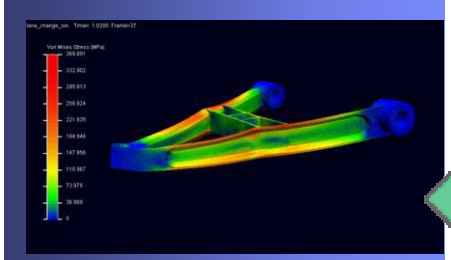

Strength

Controls & Safety

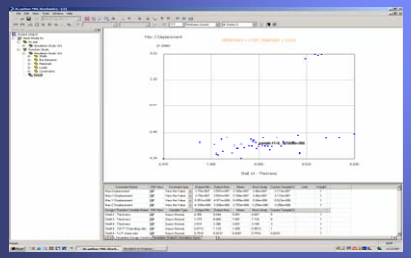

Ride/Handling

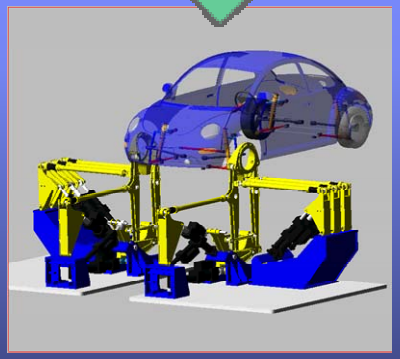

Durability

Noise/Vibration

Source: MSC.Software

Design and analysis is iterative, with multiple engineering disciplines involved in arriving at an optimum design solution. The elapsed time for each analysis and simulation task has multiple components.

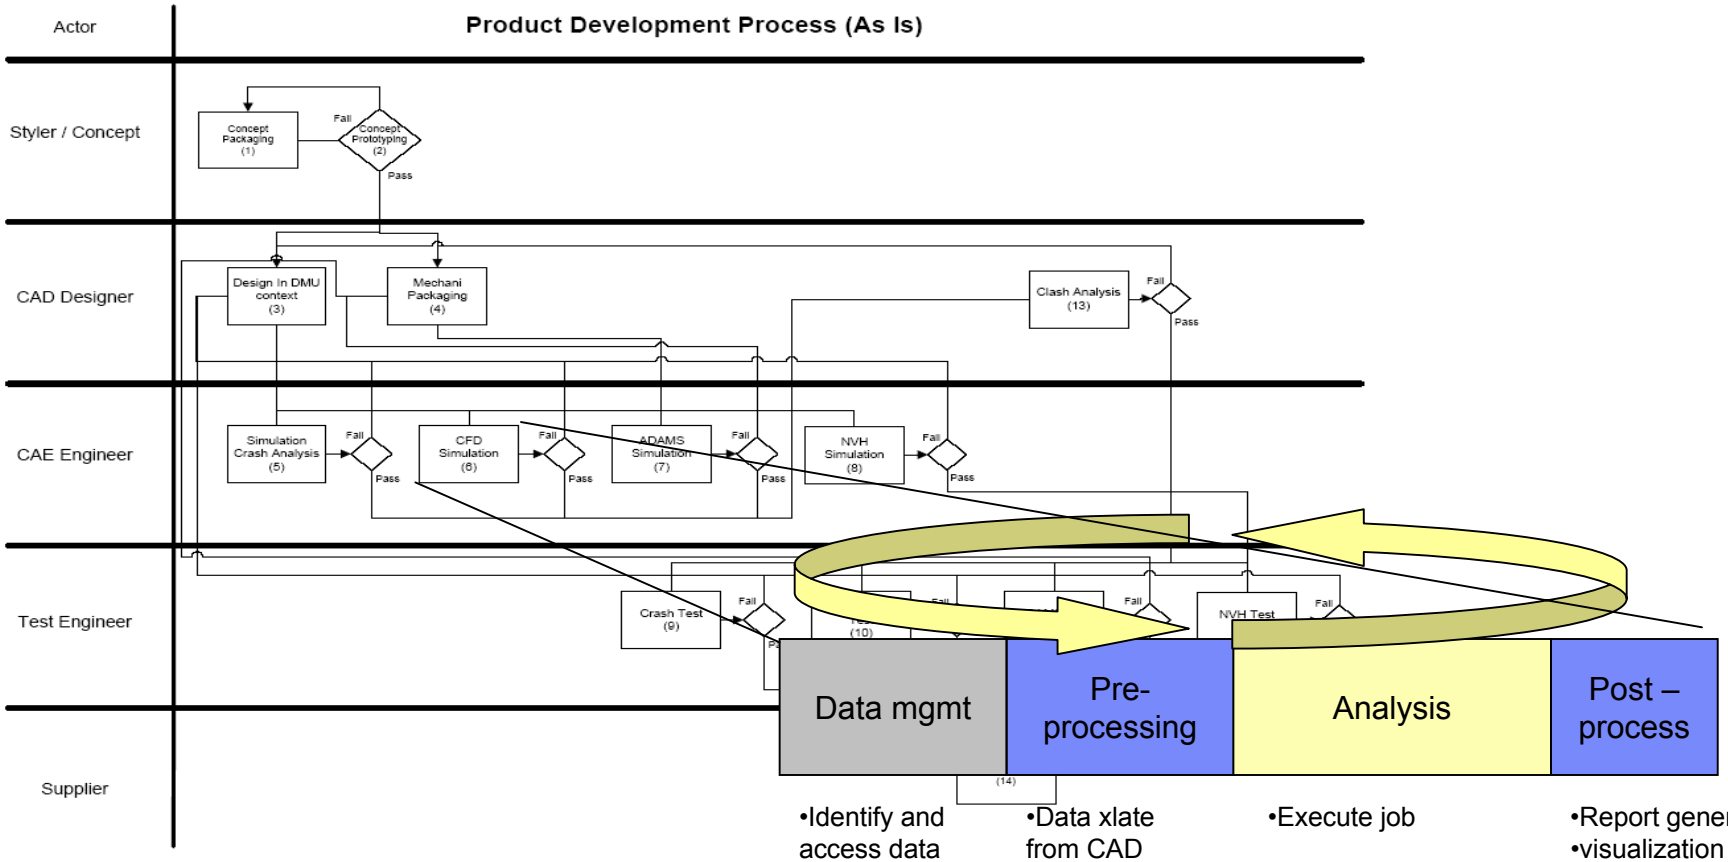

- •Identify and access data sources
- •Data reuse
- •Ensure CAD/CAE consistency
- •Manage results sets

- •Data xlate from CAD
- •Mesh creation
- •etc

- •Execute job

- •Report generation
- •visualization

Overall development cycle time is determined by the duration of each iteration and the number of iterations performed

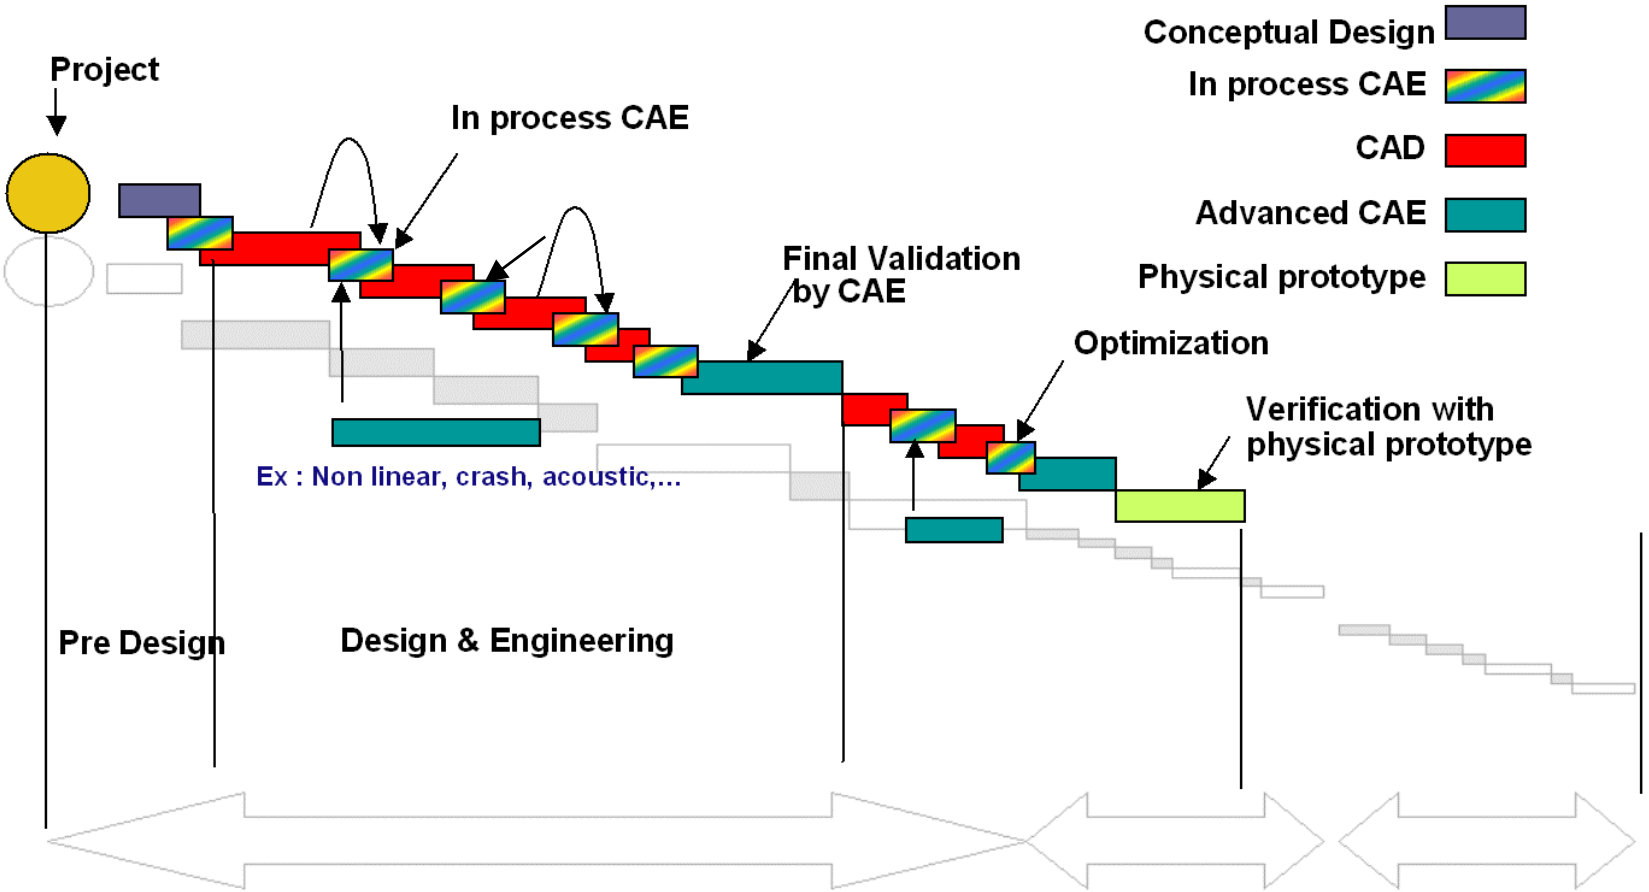

Source: MSC.Software, Dassault Systemes

Current IT infrastructure issues

- Manual, error prone, and time consuming data and storage management **consumes engineers' time (up to 50%)**
- Inability to access reliable historical data **limits reuse of product design data**
- Cannot exploit all enterprise compute resources efficiently, resulting in **overprovision in some areas and undercapacity in others**
- Need to manage resources and workload according to **business policies and priorities**
- Need to integrate design and analysis processes and data to **reduce cycle times and improve integrity of the process**
- Need to automate simulation and analysis iterations in order to **exploit full power of new simulation techniques such as stochastic analysis**
- Security issues **limit ability to share data and resources with suppliers and customers**

IBM's Automotive Engineering Innovation Framework is an integrated portfolio of offerings aiming to significantly improve the efficiency of the design / analysis cycle

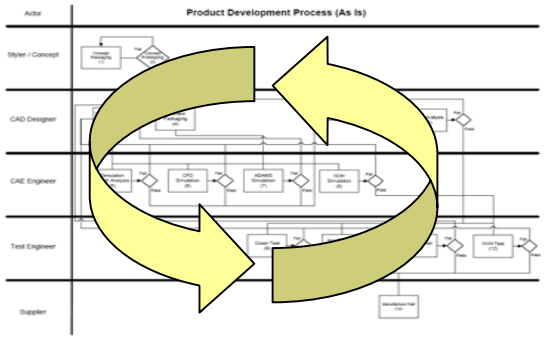

- ***Process integration & management***
- ***Design optimization***

Data mgmt

Pre-processing

Analysis

Post – process

Resource Virtualization Infrastructure

- ***Simulation Data Management***
- ***Simulation and Product Data Mgt integration***

- ***CAD /CAE integration***

- ***CAE application integration and optimization***

- ***Deep Computing visualization***

- ***Engineering Grid***
- ***Storage Management***
- ***Deep Computing Capacity On Demand***

The framework is based on a virtualized compute and data infrastructure

The AEIF architecture is based on a high performance systems platform designed to support dynamic provisioning of numerically and data intensive business processes, applications and infrastructure over the network as a service to enterprise users.

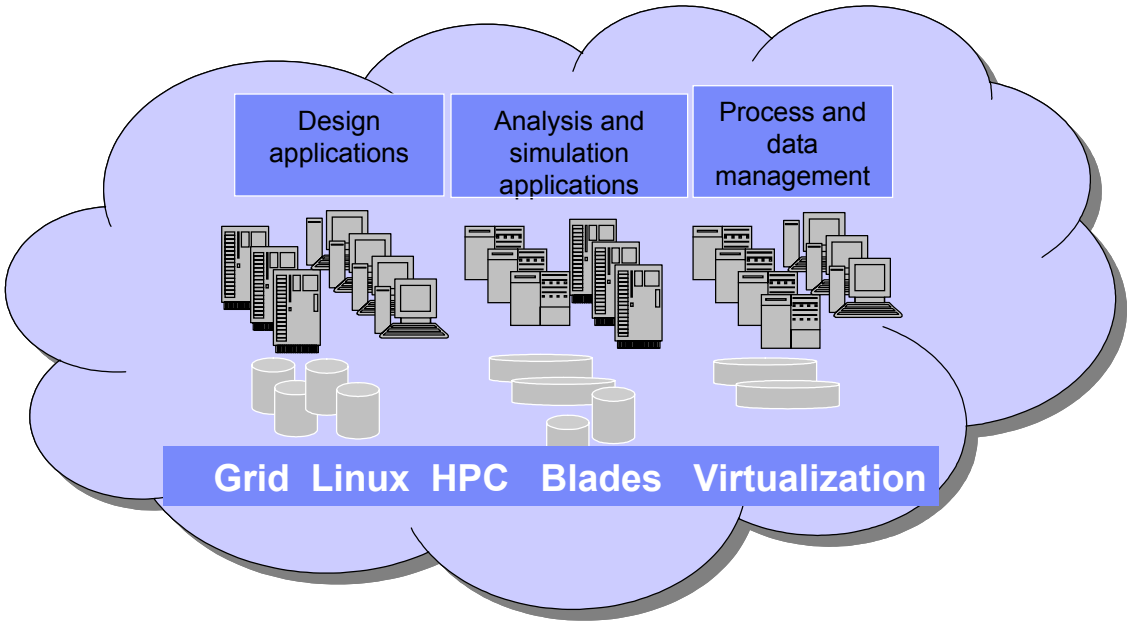

- What makes AEIF different from current approaches to infrastructure in Automotive Engineering:
  - **"Shared"** --simultaneously serves multiple users/workloads in a flexible/automated fashion.
  - **"Standardized"** --built on open standards supporting multi-vendor HW/SW integration
  - **"Scalable"**-- multi-site, multi-cluster support, with the option of additional capacity from IBM.
  - **"Service-Oriented"** -- numerically intensive computing services

New technology provides capabilities that, together, can address these demands

Core technology components

High performance 64-bit computing platforms…

Large SMP processors and DMP clusters…

High density blade servers…

Advanced data storage capabilities…

Open source operating environments….

Solution integration components

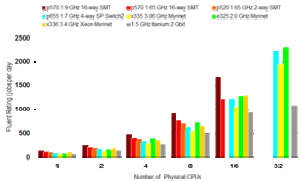

Application integration and optimization

Grid computing and resource virtualization….

'Capacity on Demand'

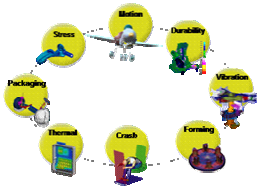

Simulation Data & Process Management / Design optimization

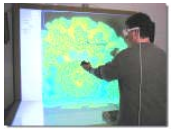

Low cost advanced visualization solutions…

Components of the Framework

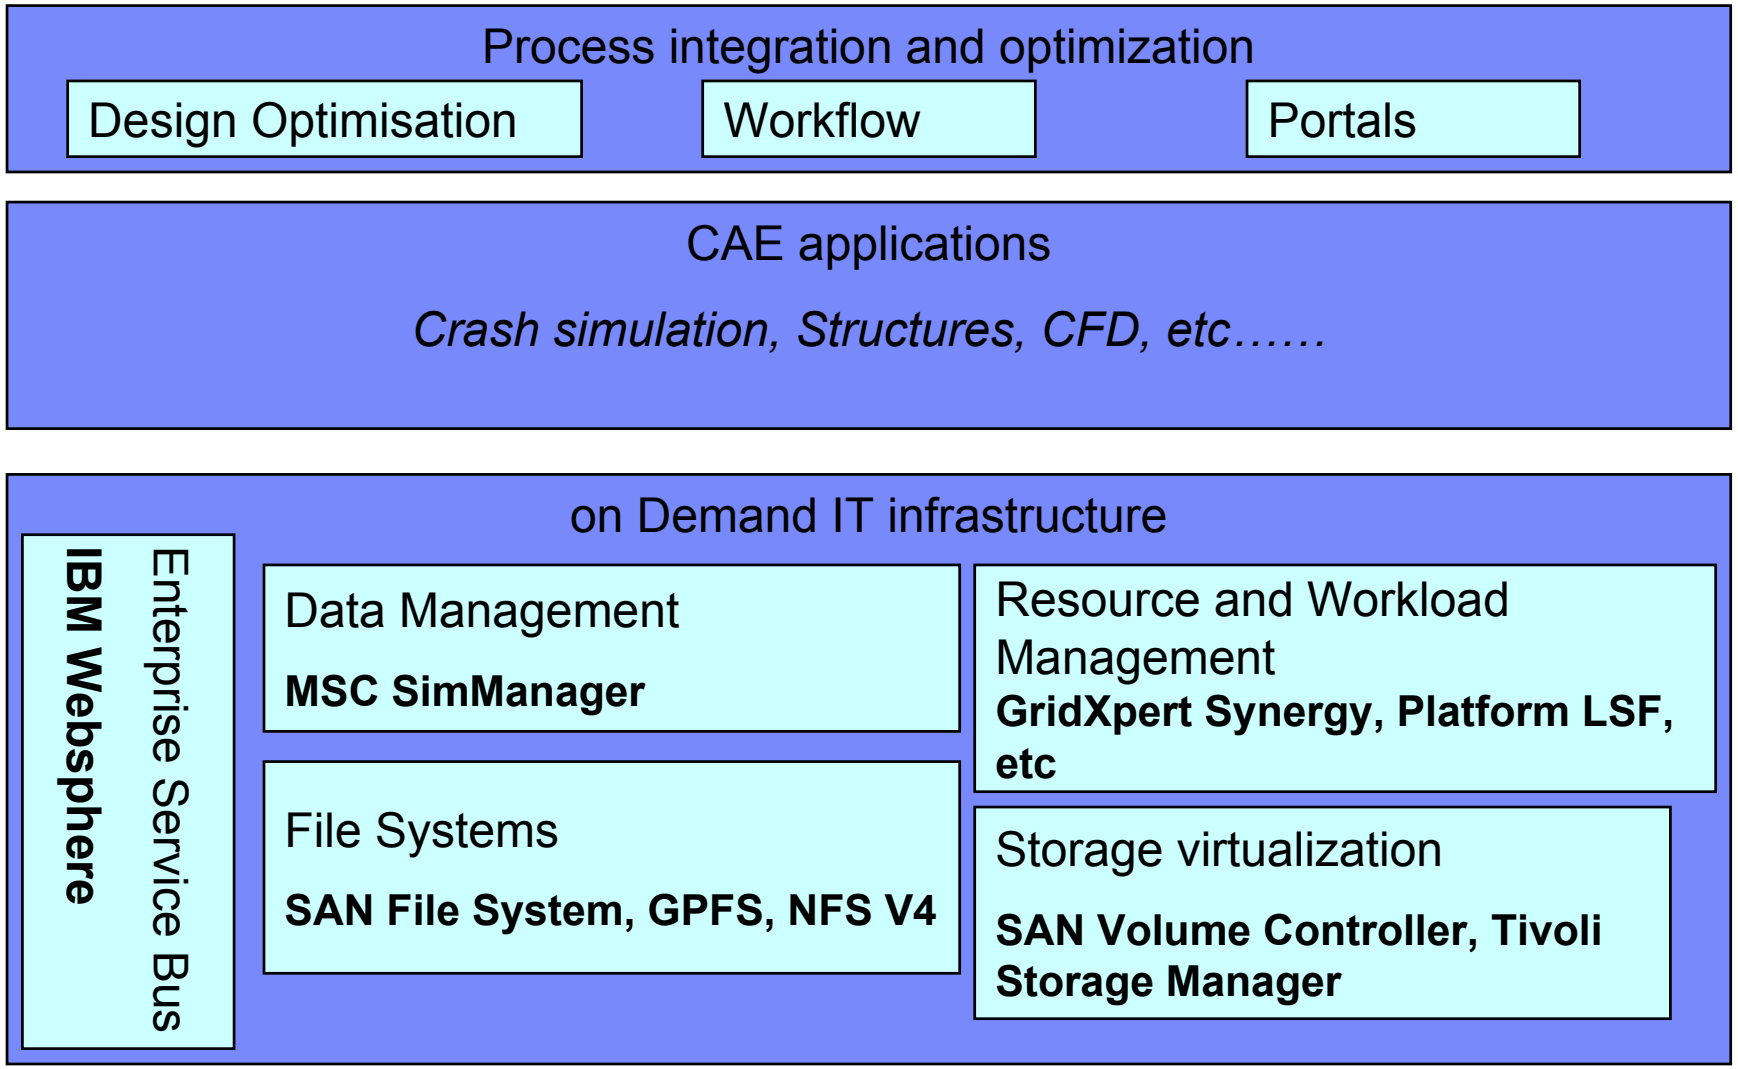

Application Integration is achieved through a standards based 'Enterprise Service Bus'

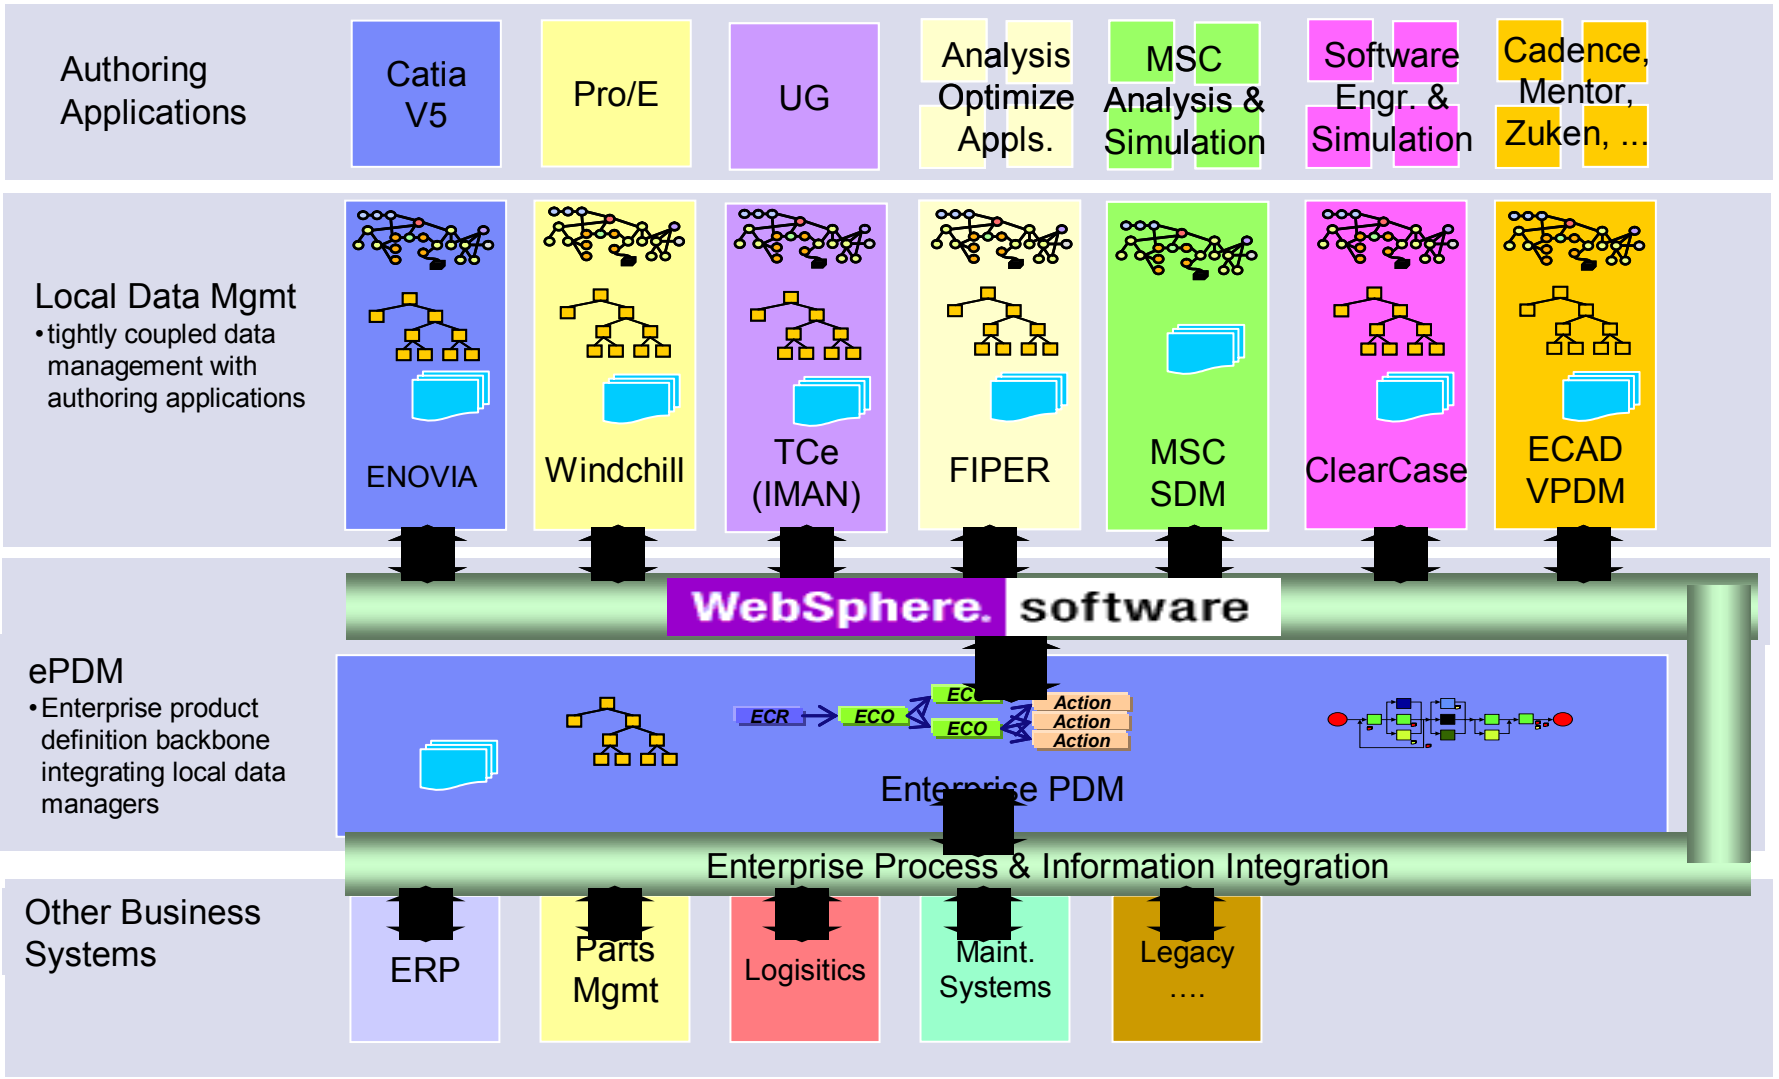

AEIF resources can be engaged in a number of ways

- **Technology briefings**
  - Help automotive customers understand the implications of the new technologies and how they can deliver value
  - Demonstrations of real world solutions to typical engineering integration and collaboration scenarios
- **Application enablement and integration**
  - Joint workshops with leading CAE ISVs to identify opportunities to exploit the AEIF and deliver added value for their customers
  - Porting, tuning and scalability testing of leading CAE codes on IBM infrastructure
- **Design workshops**
  - In depth workshops to analyse a specific customer problem, and provide a high level solution design
- **Proof of Concept projects**
  - Full realisation of solution scenarios to prove viability and reduce implementation risks
- **Implementation services**
  - Take the PoC solution and deliver into a production environment

The EMEA Centre is located at Montpellier, France

Areas of expertise…

- IBM eServer and Storage
- CAE application performance and tuning
- Deep Computing Visualisation
- Capacity on Demand
- Grid computing
- Linux cluster computing

IBM Executive Briefing Center in Montpellier

IBM

IBM @server™ and IBM TotalStorage™

The next phase of IT evolution, *e-business on demand™*, is fundamentally changing the way IT is bought, deployed and managed. An on demand business requires a technology infrastructure that is *open, integrated, virtualized* and freely enabled with *autonomic* capabilities. The IBM Executive Briefing Center in Montpellier offers highly comprehensive, in-depth technology briefings, product demonstrations and solution workshops for customers, systems integrators, independent software vendors and Business Partners wanting to gain a competitive edge. Our Briefing Center can help you understand how e-business on demand is redefining the IT operating environment, and provide you with advice on the delivery of innovative technology and solutions to meet today's ever-changing business needs.

Most Frequently Requested Topic Areas Available at the IBM Briefing Center

- Business Continuity / GDPS™
- e-business on demand™
- IBM @server™ Overview
- IBM @server iSeries™
- IBM @server pSeries™
- IBM @server xSeries™
- IBM @server zSeries™
- Server Consolidation
- IBM TotalStorage™ Solutions
- Storage Virtualization
- Linux®
- Autonomic Computing
- Grid Computing
- IBM and Oracle Solutions

Our state-of-the-art IBM Executive Briefing Center is part of the EMEA Advanced Technical Support organization, and is based at the IBM Montpellier site which also houses the pSeries and zSeries manufacturing facilities. Our 180 highly skilled professionals have received more than 5000 visitors representing 600 companies from 53 countries in 2002.

For additional topics, visit: ibm.com/servers/eserver/briefingcenter/mbc/topics.html

How you can access it…

- Briefings
- Solution Design
- Proof of Concept Projects

IBM Digital Media Solutions Targeted to Industries

Banking	*Digital Content Management* *Video Communications* *Digital Security and Surveillance* *Dynamic Digital Merchandising*
Gov	*Digital Content Management for Gov (RAMMP)* *Digital Security and Surveillance* *Digital Media Distribution*
Retail	*Dynamic Digital Merchandising* *Digital Security and Surveillance* *Digital Media Distribution*
M&E	*Digital Content Creation / Online Games* *Broadcast & Publishing Asset Mgt* *Digital Media Center*
Telco	*Marketing Asset Management* *Wireless Content Distribution* *WiFi Digital Media Gateway* *Digital Media Distribution*

RAMMP (Real-time, Adaptive, Multi-Intelligence, Multi-media Platform)

- **A platform of digital media and collaboration tools that support C4ISR processes and activities in DoD, Intelligence, Homeland Security and Law Enforcement**
  - Integrates input from multiple sensors, imaging systems and intelligence sources
  - Integrates via web services with agencies legacy systems and FSI's offerings
- **Supports ingest, processing, analysis, collaboration, archive and storage, and multi-channel bandwidth-scaled dissemination streaming**
- **Offers Live, Near-Real-Time, and Playback modes to multiple viewing devices and other Command, Control, Communication, Computing, Intelligence, Surveillance and Reconnaissance nodes.**

RAMMP Components

- RAMMP Application Families -

- **Situational Awareness** - A series of systems used to keep forces aware of the status of friendly and enemy forces over a particular locale in near real time
- **Security & Surveillance -** Systems that manage the video and sensor data and provide decision support associated with protecting life, property facilities, perimeters, borders and other critical infrastructure
- **Media Monitoring** - Data mining (archival and real time) of audio, video and print media sources
- **Forensic Content Management** - Tools to aid investigation and prosecution of crimes
- **Intelligence Analysis Support** - Tools and systems that provide support to intelligence analysts in the development of intelligence products in near real time to long range historical time frames
- **Tactical Operations Support** - Support to command and control and decision support in tactical operations such as battlefield management command and control of forces and first responders. TOS is a superset of Situational Awareness, Intelligence Analysis Support, and Security & Surveillance

- RAMMP Solution Components -

DM Core Infrastructure:
xSeries, pSeries, Blades, FastT, LTO, GPFS, Routers

Middleware:
CM, DB2/ WebSphere / MQ, Tivoli, (Digital Asset Management Suite)

IGS Assets & Services:
✓RAMMP Architecture Intellectual Property Asset
✓Installation / Integration, Consulting, Training / Education

Business Partner Applications:
NICE, Broadware, Ancept, Nexidia (FastTalk), TeleStream, ESRI, CISCO … And many others

Public Safety and Security Offering

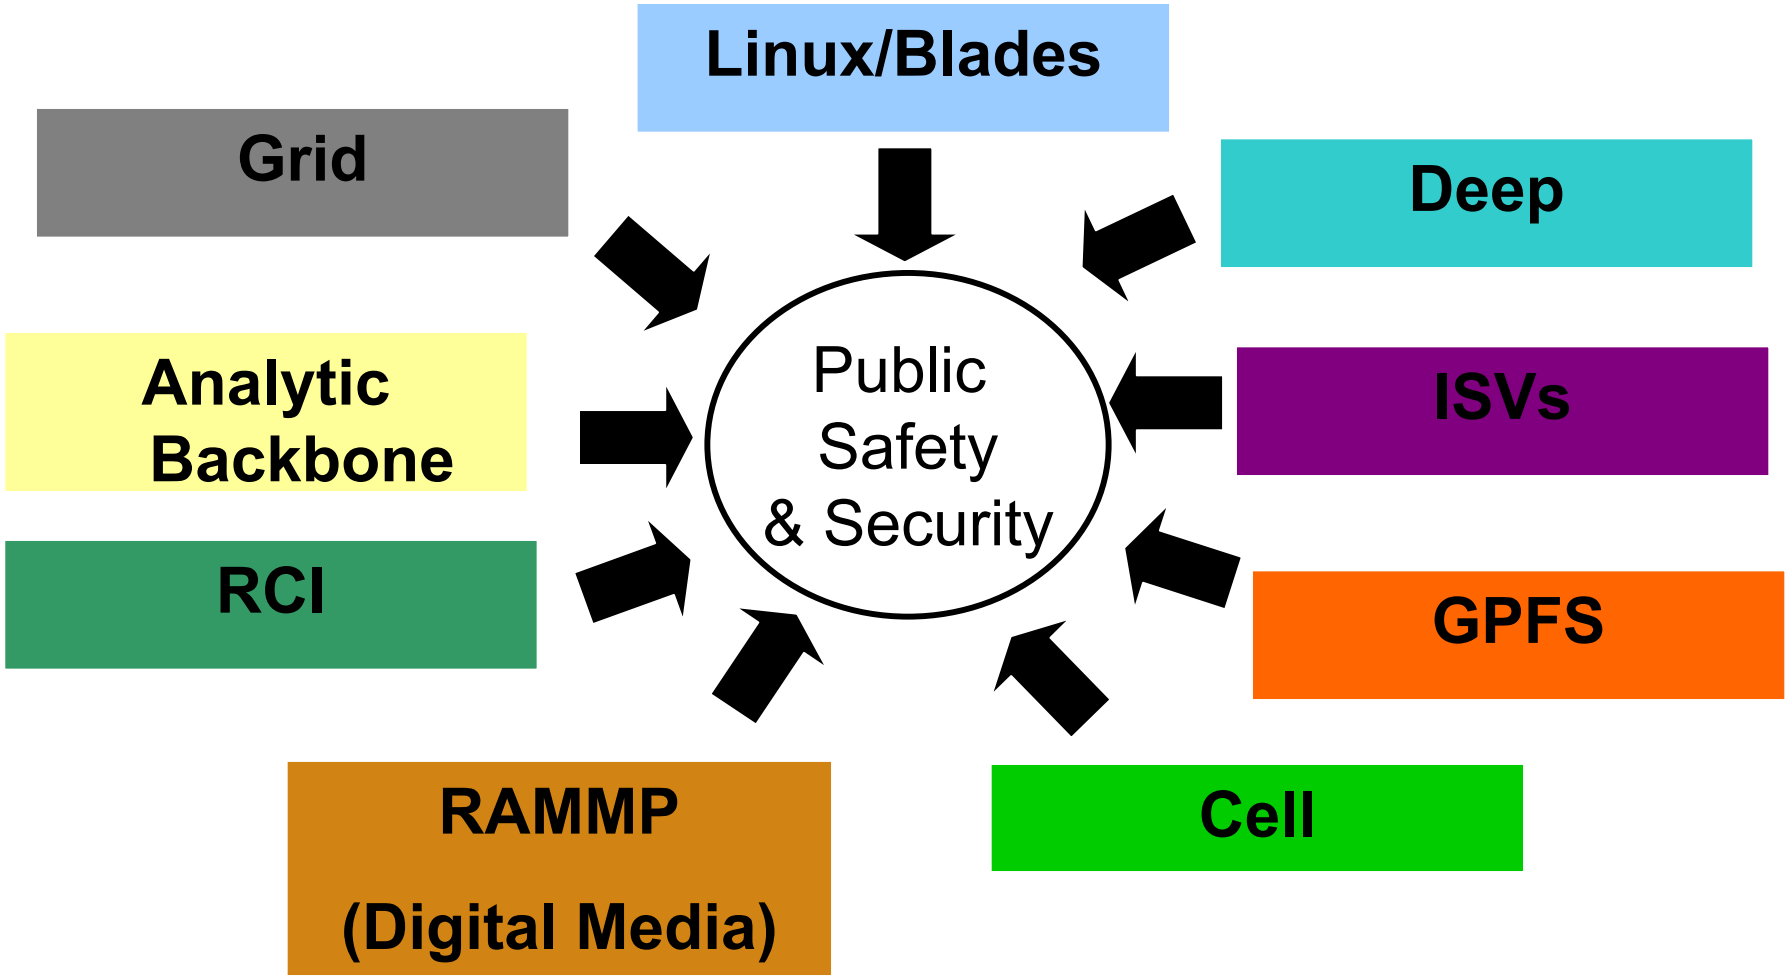

FOAK - Video Encoding for Digital Video Surveillance

- **Digital video surveillance platform based on cell processor blade**
- **Leverages H.264/MPEG-4 video compression standards**
- **Lowers cost of total ownership via reduced server and storage requirements**
- **Improves reliability via blade center design**

Digital Media Center (DMC)

- **Interconnects broadcast digital production toolsets via shared storage**
- **Eliminates production infrastructure redundancies by consolidating storage and storage management**
- **Enables production workflow simplification and effectiveness**
- **Reduces production costs by replacing proprietary elements with open standard IT elements**
- **Two offering categories:**
  - DMC Enterprise
  - DMC Transfer Managers (Data Movers)

DMC Enterprise

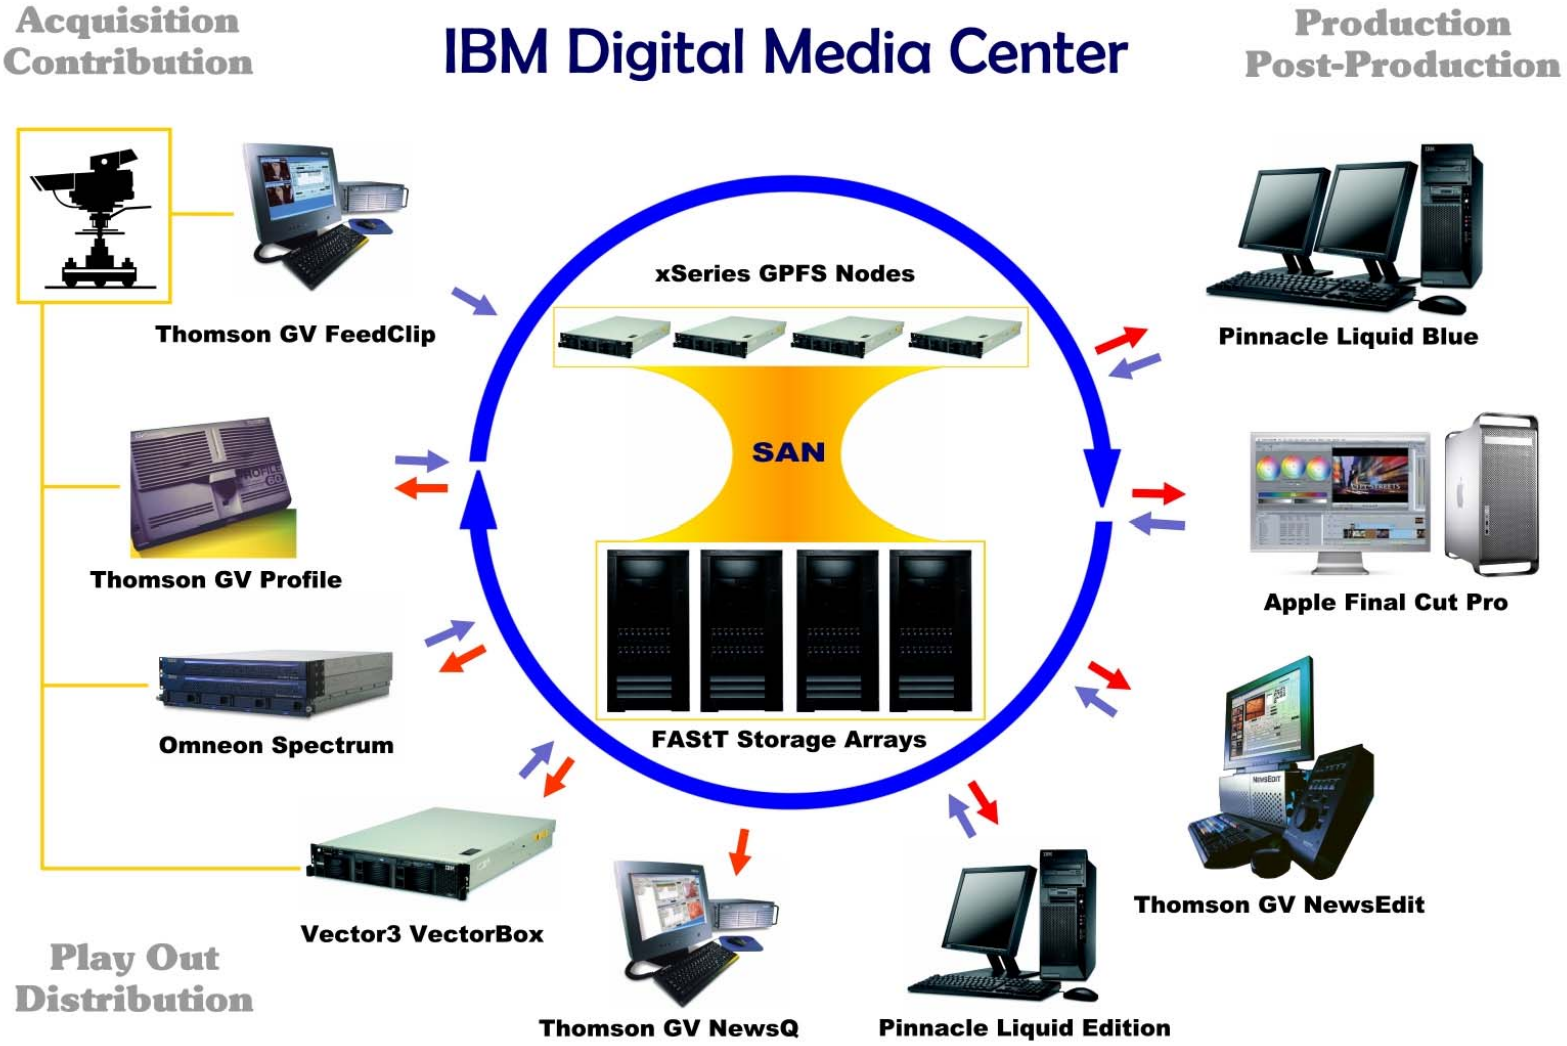

DMC Transfer Manager

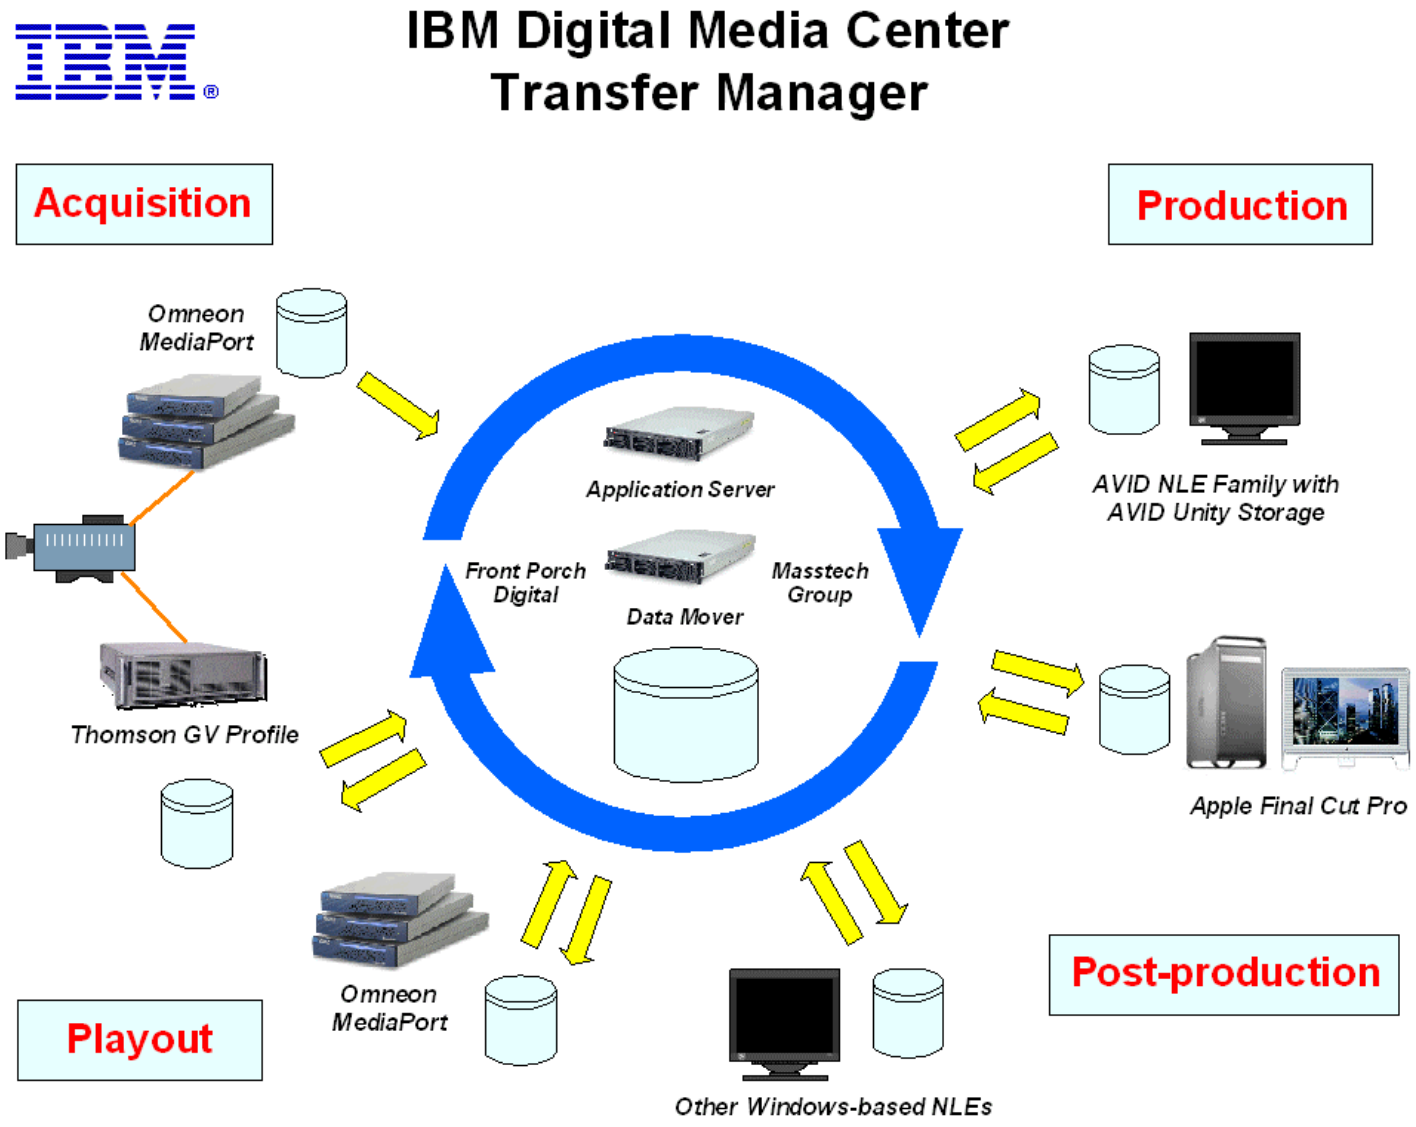

Digital Media Directions

- **On demand computing**
  - e.g., rendering on demand
- **Storage capacity**
  - e.g., TotalStorage DS4000 series
- **New server design paradigm - emphasis on cost/performance**
  - Cell processor-based systems for content creation and online gaming
  - Scalable, very high compute power at low cost
  - Extensive graphics, very high communication bandwidth

Life Sciences - Medical Archive Server (MAS)

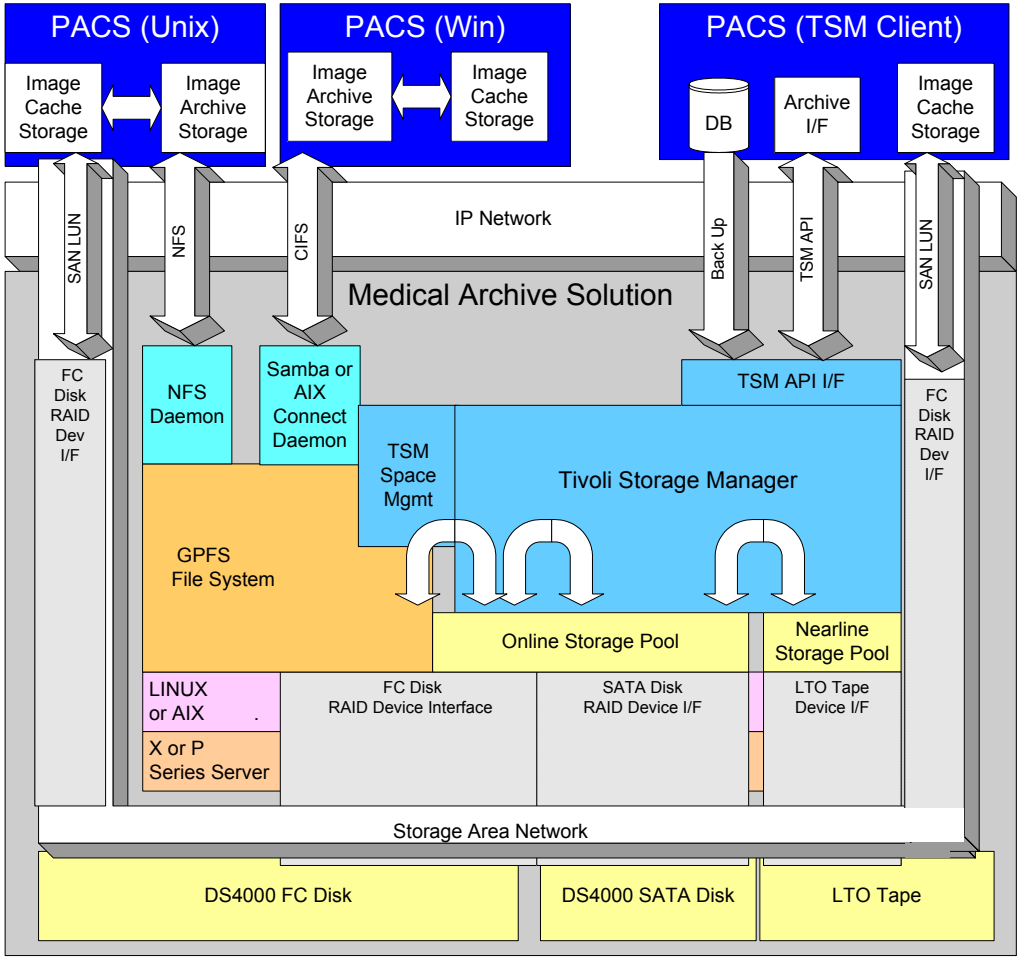

- **IBM pSeries (AIX) or xSeries (Linux)**
  - P615/P630 AIX server
  - X346 Linux
- **IBM TotalStorage**
  - DS4300 FC disk
  - DS4100 SATA disk
- **IBM LTO Technology**
  - 4560SLX
  - 3583 Automated Tape Library
  - 3584 Automated Tape Library
- **IBM Storage Software**
  - Tivoli Storage Manager
  - Tivoli Space Management
  - Fast Connect
  - GPFS (version 2)
  - DB/2 (version 2)

Healthcare & Life Sciences - Bioinformatics

Market Characteristics

- Fragmented market with many IT and application providers
- Diverse set of applications
  - Open source
  - Commercial ISVs
- Complex data environments
  - Hundreds of data formats have been identified, with new data types being defined
  - Growing exponentially & all information is saved!!
- Growing importance of real-time data analysis

Customer Profile

- Price sensitive
  - Academic/Gov't Research teams, small startups, larger biotechs
  - Buys from several vendors to meet needs
  - Writing their own code is the fastest way to get to a result / solve a business problem. Users attempt ot solve problems in the most optimal way for the given data.
  - Academic/Gov't Research customers are a "tight knit" community that shares data and collaborates
    - UCSF – big player in genetics
- The name of the game is speed
  - HPC means "high throughput"

Healthcare & Life Sciences

Dynamic Requirements for Research & More

Commodity-based computing is becoming increasingly attractive

- Users care about getting the most for their money
- Open Source Software and Linux are becoming the standard in Bioinformatics

High throughput computing

- Bioinformatics applications run "well-enough" in commodity clusters designed for maximum throughput at minimum cost

64-bit computing is emerging

- Applications can leverage features such as large memory to improve performance and solve larger research problems

IT staffs want:

- Price/performance
- Node density and space savings
- Integrated systems management tools and ease of use
- Availability features
- Ease of deployment and installation, serviceability and upgradeability

Researchers want:

- Price/performance
- Scalability
- Open source O/S and tools
- Popular applications
- 64-bit addressability
- Reliability (no down time)

Life Sciences - JS20 & OpenPower

Bring new Powerful technology to Bioinformatics/Proteomics

- **Target Market**
  - Academia/Gov't research teams, smaller startups, larger biotech
  - LE & SMB

- **Play #1: JS20 Bioinformatics Solution**
  - Applications: BLAST, FASTA, HMMER
    - Analyze sequences of DNA
  - Competitive Advantage

HMMER Performance

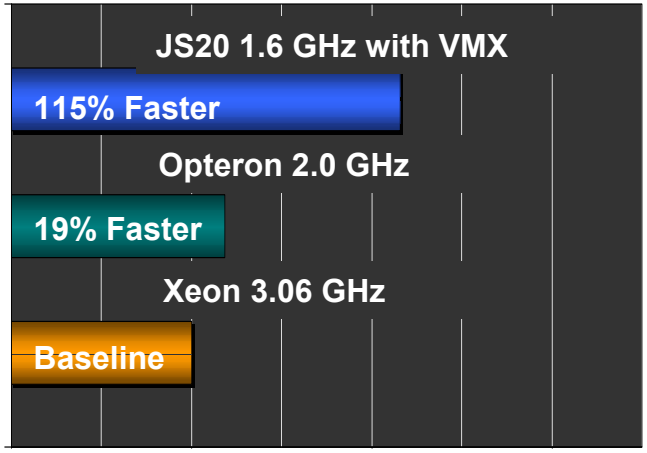

  - References: JSCC
  - More info: JS20 Sales Kit
  - Promo: University Buy one get one free

- **Play #2: JS20 Proteomics Solution**
  - Ap: Sequest 3.1 from Thermo Electron
    - Analysis & identification of proteins
  - Scale out play

Solution Stack (Recipe Page)	
Applications	**Sequest 3.1 (Thermo Electron)**
OS	**SLES9 & Red Hat Enterprise Linux V3**
Hardware	**JS20 Blade, 4MB Memory**

  - References: Medical College of Wisconsin

- **Play #3: OpenPower 720 Proteomics Solution**
  - Available late Q4 – Sell ahead
  - Application: Waters Micromass
    - Protein discovery and characterization
  - Scale Up Play

Performance of 4-way systems

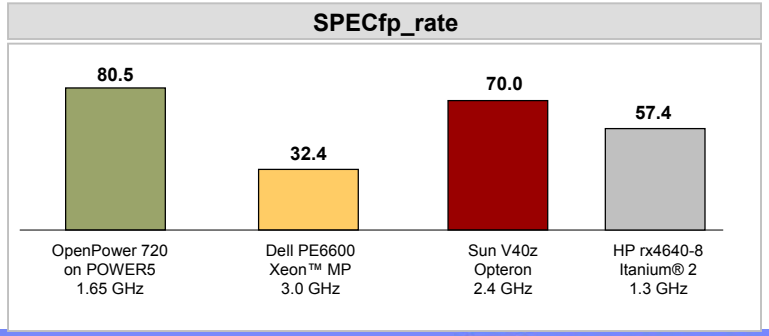

IBM eServer BladeCenter for Bioinformatics

Helping address the unique needs of pharma, biotech, academic & gov't researchers

- Affordable 64-bit computing with enhanced price/performance for key bioinformatics applications
- Scalable, designed for easy deployment and maintenance
- Available with features that can deliver high compute power on demand

Services: Implementation and Integration	
Open Source Applications: BLAST, HMMER, BLAT, FASTA, CLUSTALW, POVRAY	
OSS Middleware	Globus Toolkit
OSS Libraries	bioPERL, bioJAVA, MPICH
Compilers	IBM C/C++ and FORTRAN Compilers, GNU C/C++ Compiler
Data Management	NFS
Operating System	SUSE Linux (SLES 8)
eServer BladeCenter	NetBay Rack, 1 fully-populated BladeCenter chassis 1U Fast Ethernet switch, Expansion options

JS20 Proteomics Solution

Proteomics Solution delivers affordable, scalable performance

- **Marketplace**
  - Proteomics uses computational tools in the study of proteins.
  - Large opportunity for 64-bit optimized performance
  - Market is moving to Linux
  - IBM has a great relationship with the top ISVs
- **Value proposition**
  - JS20 and IBM TotalStorage will allow Proteomics customers to process and store more data proving a quicker time to solution at a price point that is competitive with Intel Architecture based systems which will maximize customer time-to-value
- **Selling Tools**
  - Comprehensive Internet solution external web site
  - IBM eServer BladeCenter JS20 White-Paper
  - JS20 solution for Proteomics Brochure
  - Performance proof-point: Active Protein DB search over 40x faster…Thermo Electron
  - Customer reference: **Medical College of Wisconsin**

Sector: Public
Focus area: Industry app in LE & SMB
Industries: Pharma, Gov't, academia
Size: $12M LS 2004 GA: 09/30/04

Solution Stack	
Applications	**Sequest 3.1 (Thermo Electron)**
Middleware & Tools	**Apache, Java, OTG diskXtender (optional)**
System Mgmt	**TSM (optional)**
OS	**SLES9 (first) & Red Hat**
Hardware	**JS20 Blade, TotalStorage Cisco GB Enet Switch**

Instruments

I/O Intensive

SEQUEST Cluster Application Server

CPU, Memory & I/0 Intensive

SEQUEST Cluster Nodes

CPU & Memory Intensive

Mare Nostrum

- **#1 Supercomputer in Europe, #4 Worldwide (Top500 11/04)**
- **Customer: Ministerio de Ciencia y Tecnologia**
  - Barcelona, Spain
- **2282 IBM eServer BladeCenter JS20 2.2 GHz Blades**
  - All commercial components
- **Linpack performance @ 1782 Nodes**
  - Rmax 20.53 teraflops
  - Rpeak 31.363 teraflops
- **Metrics**
  - Speed/Power: 50 GFLOPs*/Kwatt
  - Speed/Space: 13 GFLOPs*/Sq Ft
  - Price/Speed: $970K/TFLOP**
- **Applications include**
  - Life Sciences: proteomics, bioinformatics, computational chemistry
  - Weather
  - Material science

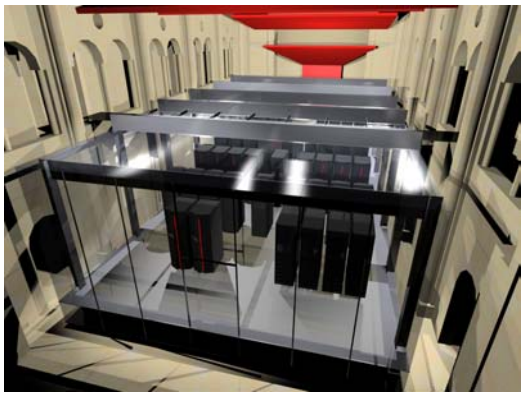

* Rpeak
** Rmax

Business Intelligence

Accountability and Compliance

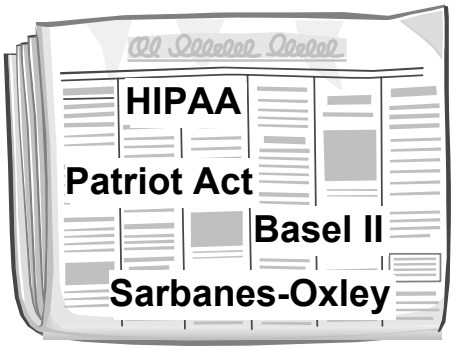

- Risk Management
- Fraud and Abuse
- Public Protection

Customer Knowledge

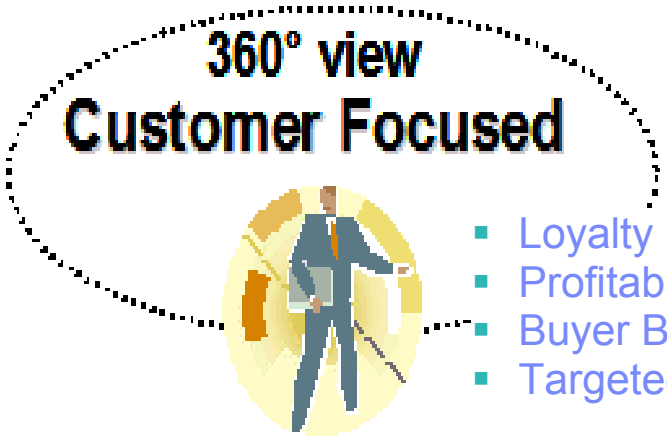

- Loyalty
- Profitability
- Buyer Behavior
- Targeted Offers

Homeland Security

- Internet Buzz
- Anti-Money Laundering
- Border Control
- Crime Information

Business Performance

- Globalization
- Business Controls
- Mergers and Acquisitions
- Supply Chain Efficiencies

Business Intelligence - Massive Data Requirements

- **A 2003 report* on extract, transform and load (ETL) issued an 18 month trend forecast including:**
  - Move towards near real-time will account for 19% (3x growth) of all DW
  - Average DW will be processing 100-150GB per hour
  - Largest DW will exceed 50TB
  - Average extraction from 12 distinct sources
  - 86% of the respondents rated reliability as “very important” with a statement of “users want greater reliability above all else”

** The Data Warehousing Institute (TDWI), report based on 756 respondents*

Leadership eServer Performance

- **100GB TPC-H Performance Leadership**
  - #1 IBM eServer 325/DB2 UDB 8.1 (12,216 QphH)
  - #2 IBM eServer OpenPower 720/DB2 UDB 8.2 (6,357 QphH)
- **300GB TPC-H Performance Leadership**
  - #1 IBM eServer 325/DB2 UDB 8.1 (13,194 QphH)
- **1 TB TPC-H Price/Performance Leadership**
  - IBM eServer p5 570/DB2 UDB 8.2 (53.00 US $ Price/QphH-26,156 QphH)
- **TPC-C Performance Leadership**
  - #1 IBM eServer p5 595/DB2 UDB 8.2 (3,210,540 tpmC)

Balanced Configuration Unit (BCU)

Minimum replicable HW/SW stack necessary to start or expand the infrastructure of a Business Intelligence system

- ✓ Reusable and validated approach to data warehousing
- ✓ Improves sizing and capacity planning
- ✓ Highest levels of performance and scalability
- ✓ Reduces risk, time to business value and total cost of ownership (TCO)
- ✓ Provides scalable performance ratio of disk I/O to memory to CPU to network

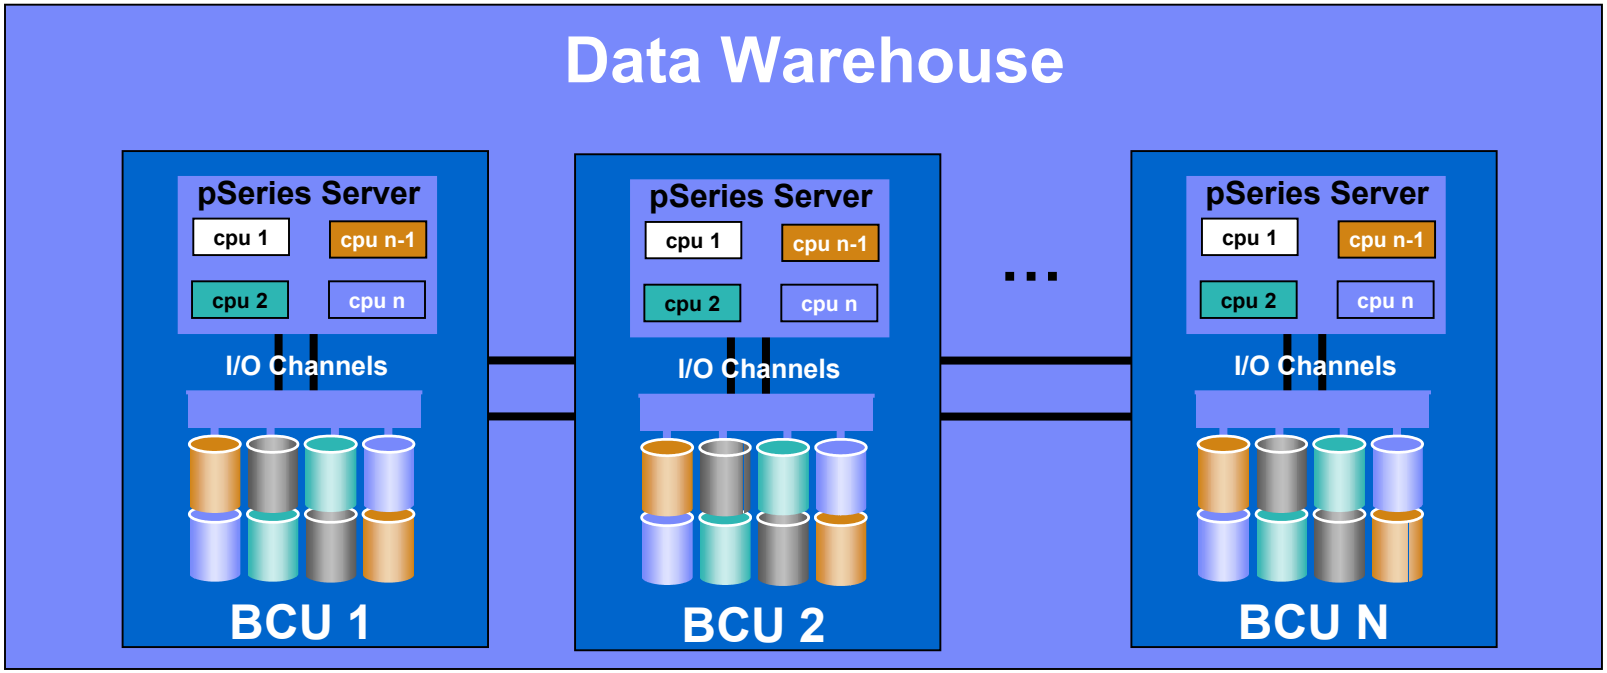

BCU-based Data Warehouse Example - p655 & DS4500

- **Flexible warehouse architecture to accommodate mixed workloads, consolidation of data marts, expand user base**
  - Increase availability with recoverable system and database
  - Simplify systems, database and architecture management

Servers	**2 p655 (8-way, 1.7GHz with 32GB memory)**
Storage	**2 DS4500 with 2.5TB each**
Additional Notes	**High Availability design** **8 logical database partitions** **- 200GB per LDBP (15 disks per)**

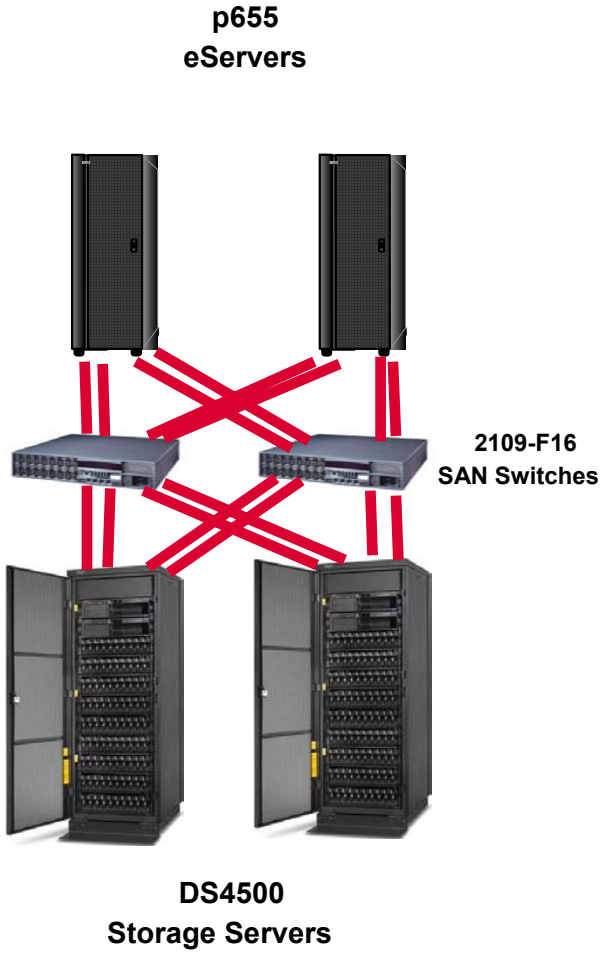

BCU Portfolio (Servers)

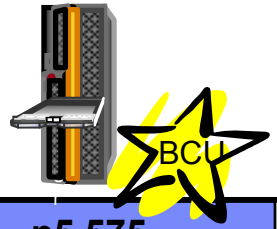

	p5 570 (Fast Track)	p5 575 (Long-term)	e326	x346
Operating System	UNIX (AIX 5L)	UNIX (AIX 5L)	Linux	Linux
Processor Technology	POWER5	POWER5	AMD Opteron™	Intel® Xeon™ (support for em64t)
Processor Speed	1.9 GHz	1.9 GHz	2.4GHz	3.6GHz
# of Processors	8-way for BCU*	8-way	2-way	2-way
Memory	Up to 256GB*	Up to 256GB	Up to 16GB	Up to 16GB
Form Factor	8U rack*	24” system frame nodes	1U rack	2U rack
RBI Rating		est. 10.0		

High performance compute node with ultra-dense clustered packaging design

**scalable to 16-way with up to 512GB memory and 16U*

IBM TotalStorage® DS4500

Superior price/performance for data warehouse storage

- **High Performance 2Gb Fibre Channel Connectivity**
  - Up to 148,000 I/O's per second, 772 MB/sec
- **Scalable Rack-Mounted Storage**
  - 36 GB to 32.8 TB raw disk capacity (56 TB with SATA)
  - Supports up to 224 hot-pluggable disks;
    - 36, 73, 146GB 10K RPM; 18, 36, 73GB 15K RPM; 250GB 7200 RPM
- **High Availability**
  - RAID levels 0,1,3,5,1+0 with global hot spares
  - Dual RAID controllers, redundant paths with failover protection
  - Redundant power and cooling; dual line cords
  - 2 GB mirrored battery protected cache
  - FlashCopy, Volume Copy and Enhanced Remote Mirror
- **DS4000 Storage Manager**
  - Powerful, intuitive storage managment tool
  - Partition storage for pooling
  - Dynamic “Anytime Administration” and recovery guru
  - Service Alert “call home” facility
- **3-year Warranty***

Storage Performance Council
SPC Benchmark-1™ Certified
www.storageperformance.org

BCU disk of choice!

"SAN-ready"

High Performance, Scalable Storage for AIX, Windows NT/2000, HP-UX, SUN, Netware, Linux

Storage Design Comparisons

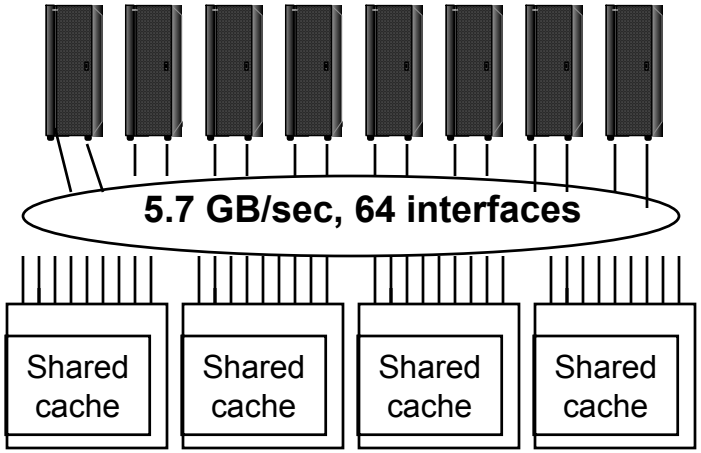

SAN

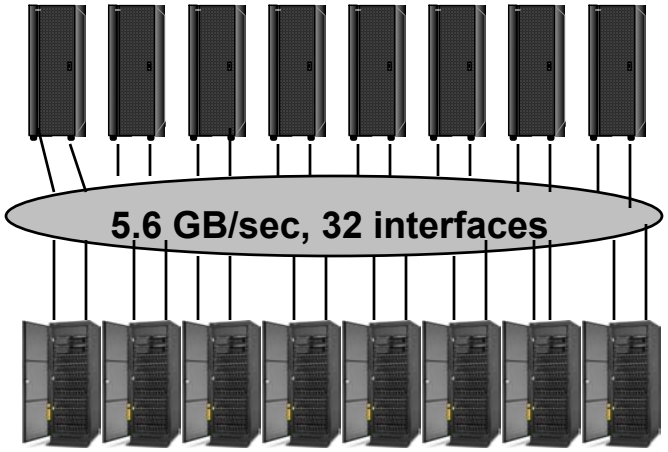

Up to 50% less square footage

EMC or Hitachi Monolithic Design

- **Violates share nothing architecture**
- **Cache contention**
- **Limited scalability, vertical growth**
- **Hard to predict overall performance increase**
  - No BCU capability
- **High Price / Performance**

IBM DS Modular Design

- **Preserves share nothing architecture**
- **Limitless scalability**
- **Take advantage of most recent technology**
- **Predictable overall performance increase (BCU)**
- **Low Price / Performance**

Storage Virtualization Example

- The Company:
  - SC Johnson Wax (3Q04 win - large manufacturer of retail consumer package goods)
- Business Needs:
  - Infrastructure for SAP Business Warehouse
- IBM Solution:
  - 4 p5-550, 1 p5-570 eServer servers
  - 9 DS4500 TotalStorage systems, total 55 TB physical disk capacity
  - 2 four-node SVC clusters for synchronous Remote Copy
- "Sealed-the-Deal":
  - BCU concept in SAP BW design
  - IBM storage virtualization strategy

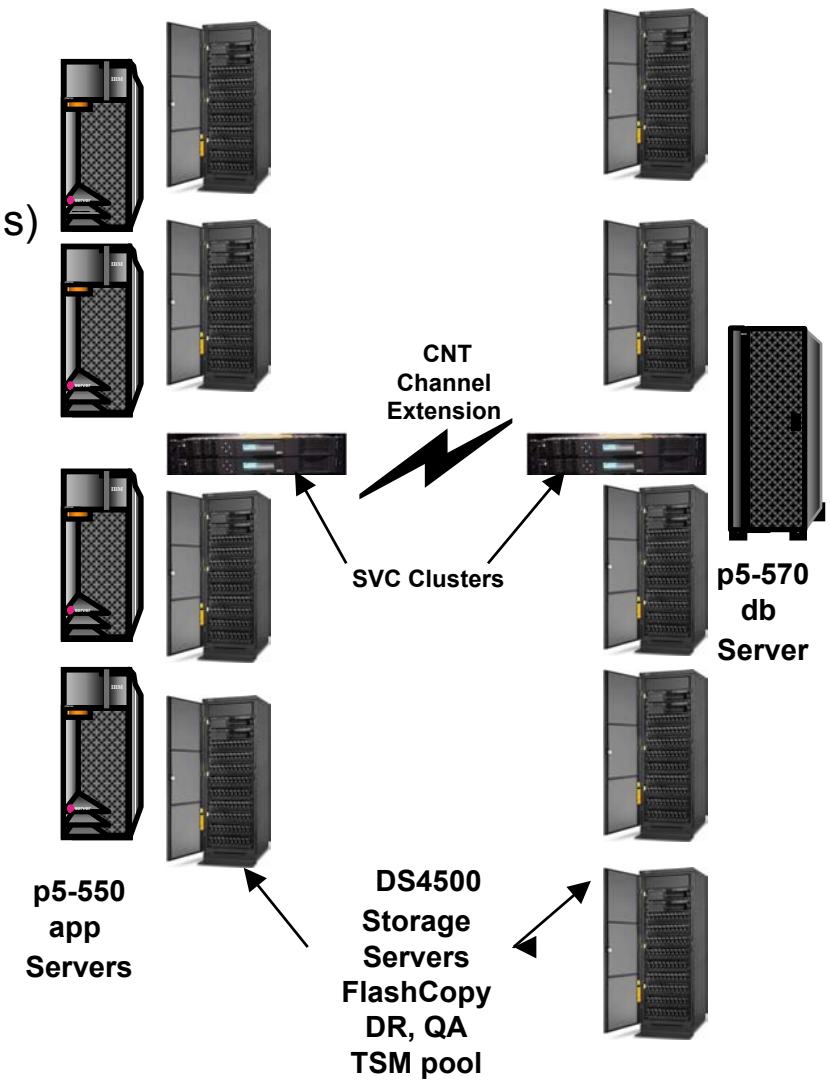

IBM BI Solutions For eServer & TotalStorage

Data Analytics with Strategic Alliance Partners

- PeopleSoft
- Siebel Analytics
- SAP Business Warehouse (BW)
- SAS

Database With Built-in BI Features

DB2 Information Management Software

- DB2 Universal Database (UDB) Data Warehouse Editions (DWE) and BCU

Infrastructure Opportunity

Custom Infrastructure

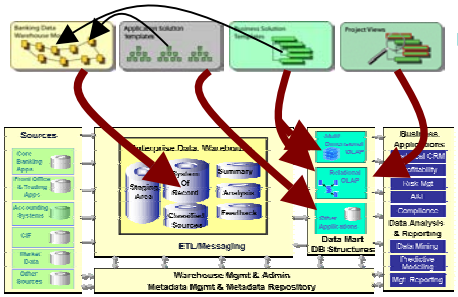

- IBM Information FrameWork (IFW) & Banking Data Warehouse (BDW)

Data Mart Consolidation With IBM eServer & TotalStorage

- ✓ Cost savings
- ✓ Flexibility
- ✓ Risk reduction

Opportunities with ISVs (PeopleSoft, SAP, Siebel, etc…)

Conclusion

- **Deep Computing systems are crucial and will continue to drive advancement in the computer industry**
- **Petaflop performance will advance in response to the needs of the scientific community**
- **Growing application complexity demands adaptive Deep Computing systems**
- **On demand models of computing are meeting the needs of an increasing and diverse segment of the marketplace**

Selling Deep Computing

- **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, …
- **Know what to sell … seek help**
  - One size does not fit all
  - Broad product and solution portfolio – HW, SW, services, applications/tools
  - 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.g., ACTC, Computational Biology Center)
  - Pricing and profit (e.g., pSeries LSRB)
- **Contact your WW or geo Deep Computing sales experts**
  - Deep industry and technology expertise
  - Brand specialists (pSeries, xSeries/Linux clusters, storage)

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

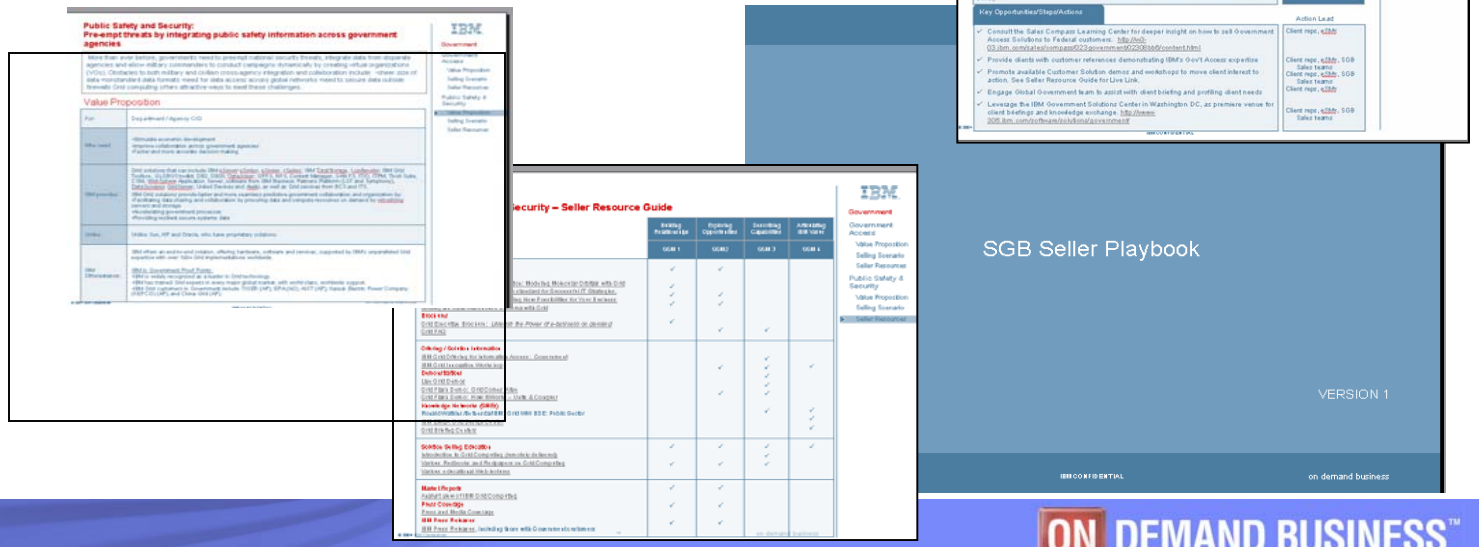

Additional Resources

- 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

End Of Presentation

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