IBM

Innovation & Technologies for the on demand world

Michel Teyssedre
Vice President – EMEA Strategic Business Opportunities

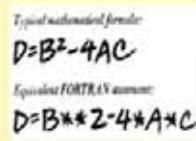# History of Innovations

1944: Mark 1   1948: SSEC   1956: RAMAC   1957: FORTRAN   1966: One-Device Memory Cell   1967: Fractals   1970: Relational Database

1971: Speech Recognition   1973: Winchester Disk   1979: Thin Film Recording Heads   1980: RISC   Nobel Prizes   1994: SiGe   1993: RS/6000 SP 1996,97: Deep Blue
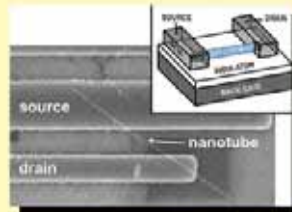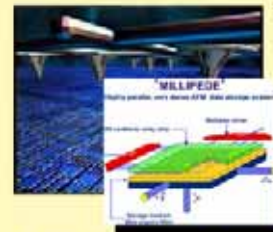
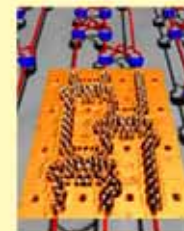1997: Copper Interconnect Wiring   1998: Silicon-on-Insulator   1998: Microdrive   2001: Nanotube Transistor   2002: Millipede   2002: Molecule Cascade Logic Circuit
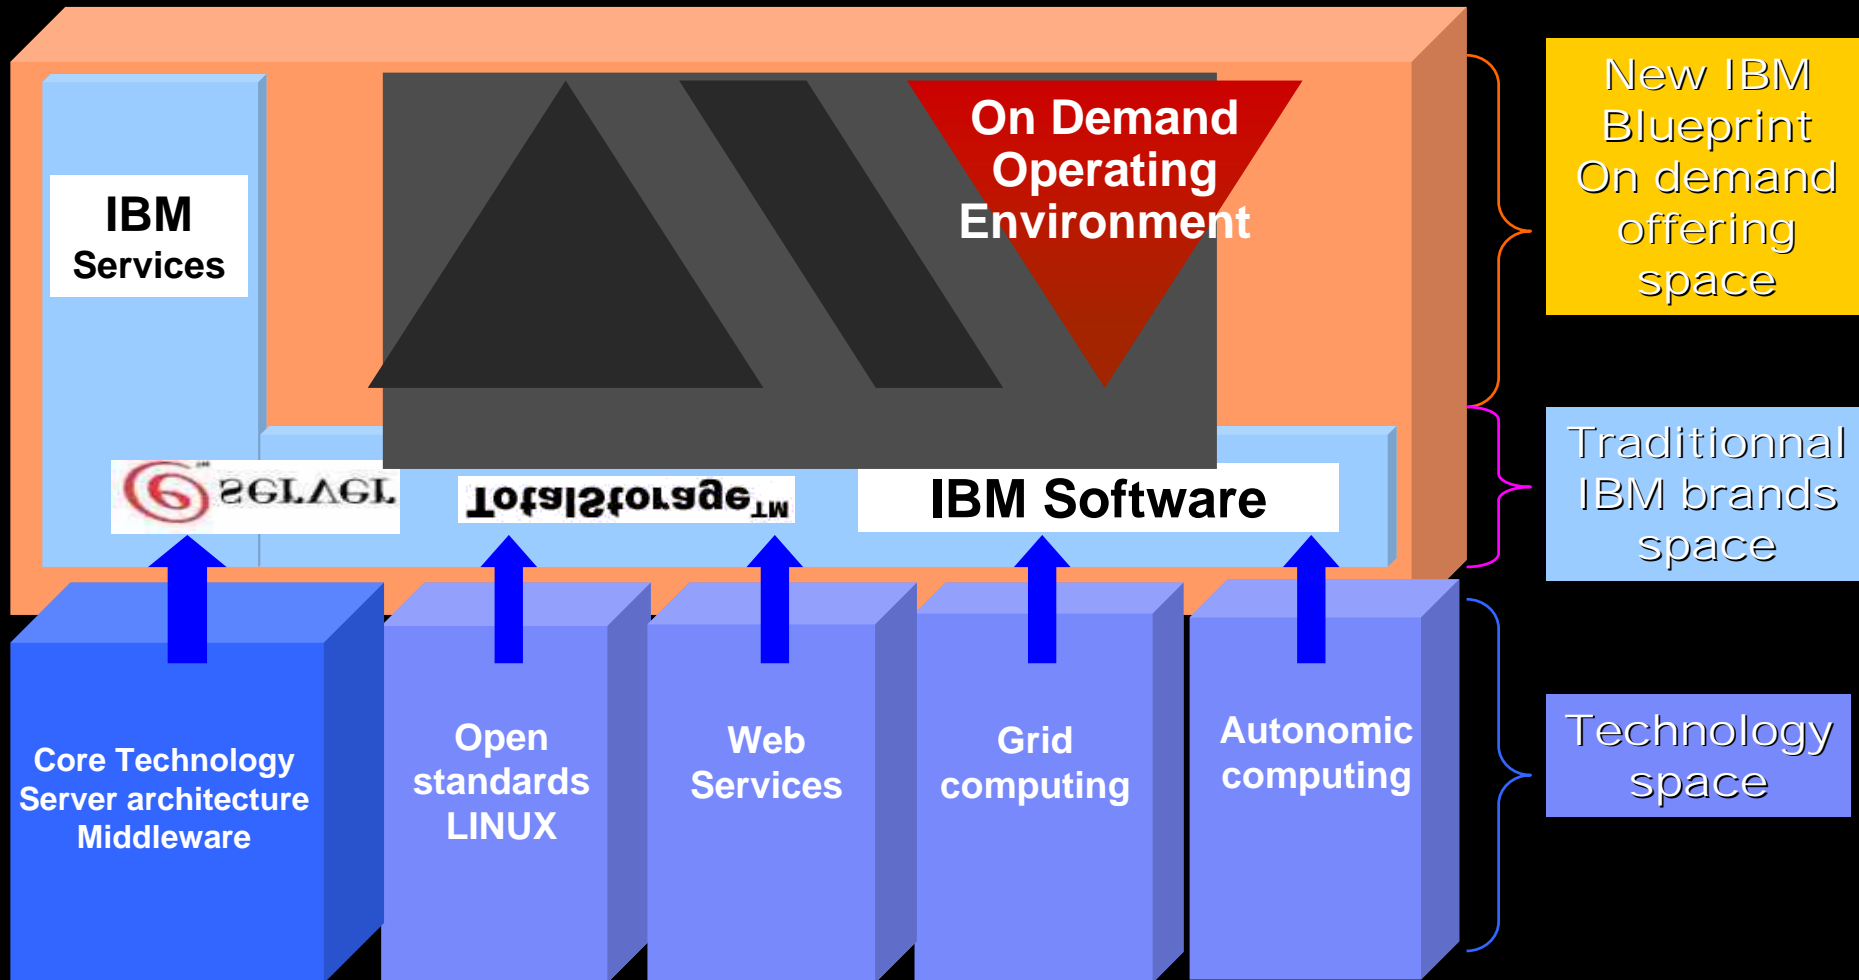
Technology Leadership
POWER EVERYWHERE

Servers

POWER5

POWER4+

2001

POWER4

2000

2002

POWER3

1998

IBM

PPC 75

Nanotechnology

1997

POWER2

Millipede

PPC 750CXe
low- k dielectric

SiGe

Strained Silicon

PPC 440GX

BLUE GENE

Copper Circuitry

PPC 603e

Silicon-On-Insulator

PPC 440

PPC 405G 2

X-Architecture

TotalStorage

PP

z900 multichip module

POWER IV

… and we deliver on our commitments

New tools for the new markets

Flexible server growth

Scale UP

Scale Up / Consolidated Platforms

Clusters
- Linux
- HPC
- AIX 5L
- Parallel Sysplex

i890/p690/z900
- Logical Partitioning
- Virtual Servers
- Server Consolidation

Enterprise SMP

x440
- Scalable Nodes
- Partitionable
- Modular growth

High-Density Rack-Optimized

BladeCenter
- Rapid Deployment
- Flexible Architectures
- Switch Integration

Blue Gene

Chip (2 proc)
2.8/5.6 GF/s 4MB

Compute Card (2 chips, 1x2x1)
5.6/11.2 GF/s 0.5 GB DDR

Node Board (32 chips, 4x4x2) 16 Compute Cards
90/180 GF/s 8GB DDR

Cabinet (32 Node boards, 8x8x16)
2.9/5.7 TF/s 256 GB DDR

System (64 cabinets, 64x32x32)
180/360 TF/s 16 TB DDR

Scale OUT

Scale Out / Distributed Platforms

Supercomputing Roadmap
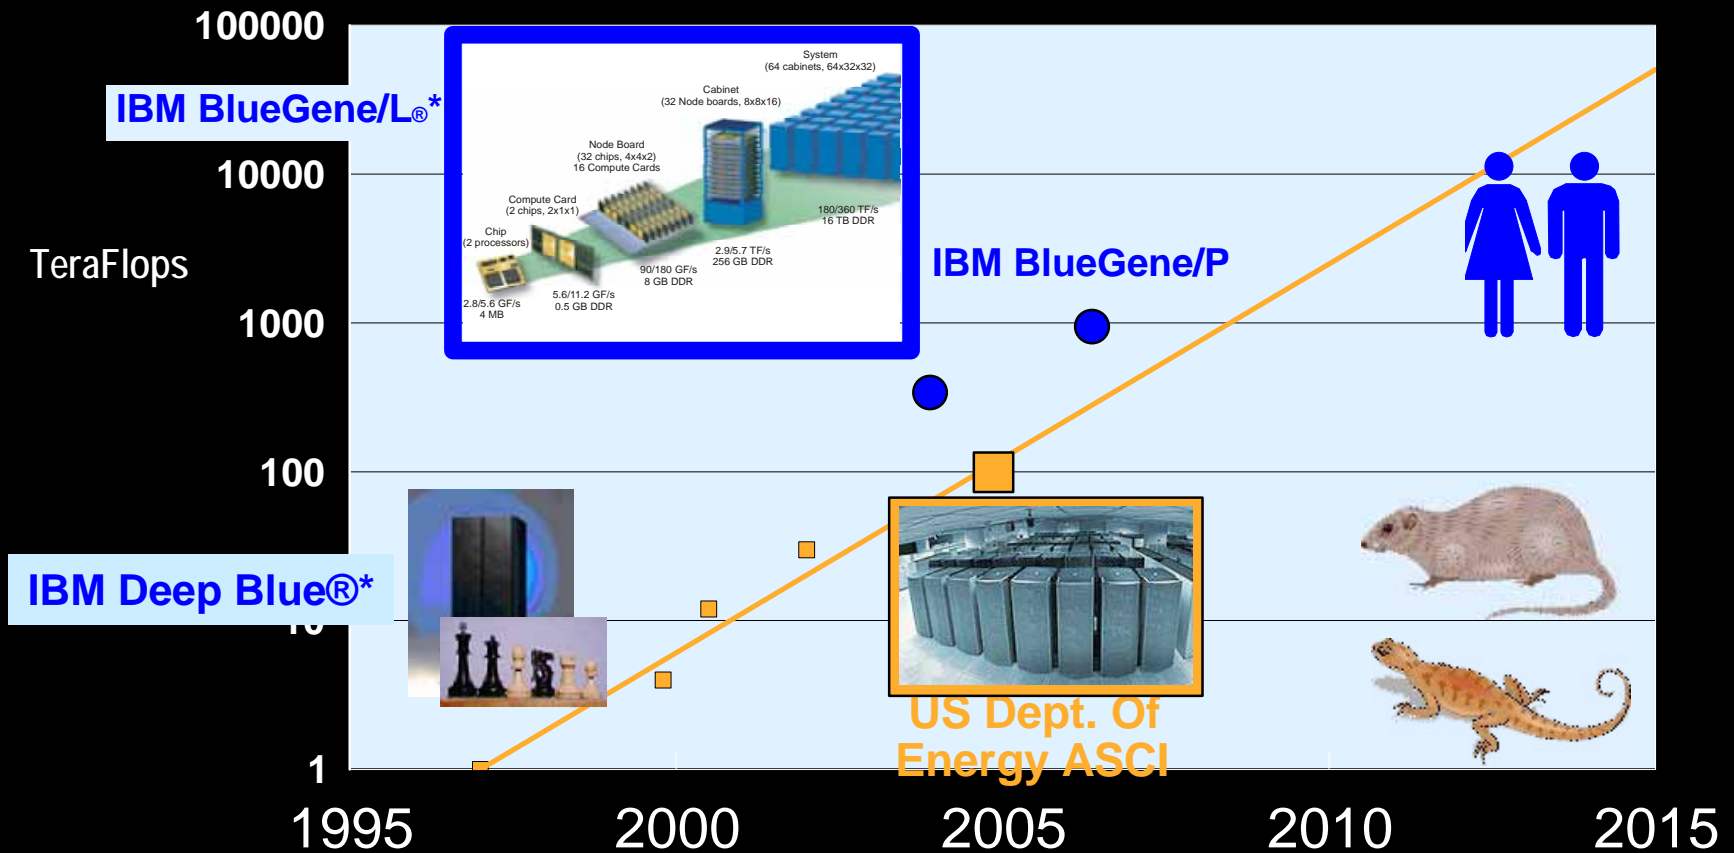
Source: ASCI Roadmap www.llnl.gov/asci, IBM

Brain ops/sec: Kurzweil 1999, The Ace of Spiritual Machines

Moravec 1998, www.transhumanist.com/volume1/moravec.htm

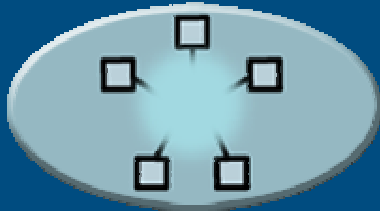**Open standards LINUX**

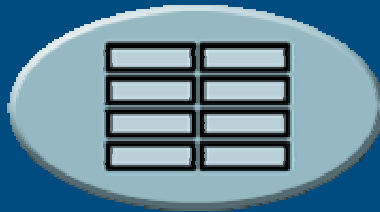**Grid computing**

Autonomic computing

Virtualization Engine
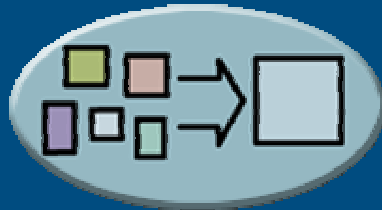
Strategy: Focus Areas

Continued Focus

Infrastructure Solutions

Linux Clusters/ Blades

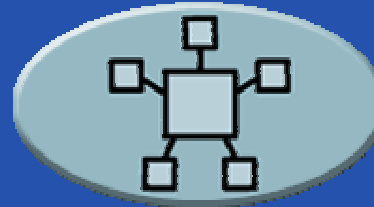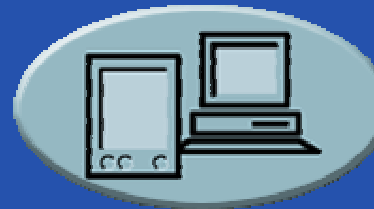Workload Consolidation

Industry Applications

New Focus

Business Partner Value Networks

Emerging Countries

Linux Client

IBM e-business on demand operating environment blueprint

Grid computing
Develop technologies for generic resource sharing

Grid computing allows integration across the enterprise

I/O   **Storage**   **Operating System**

Processing   **Applications**   **Data**

Grid Computing :
Distributed Computing Over a Network, Using Open Standards to Enable Heterogeneous Operations
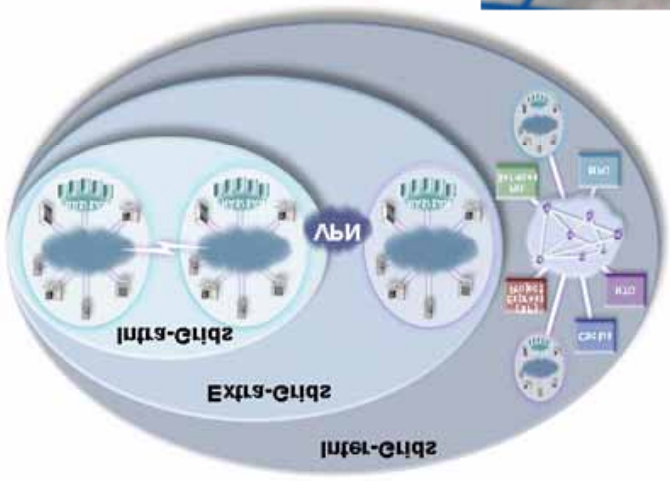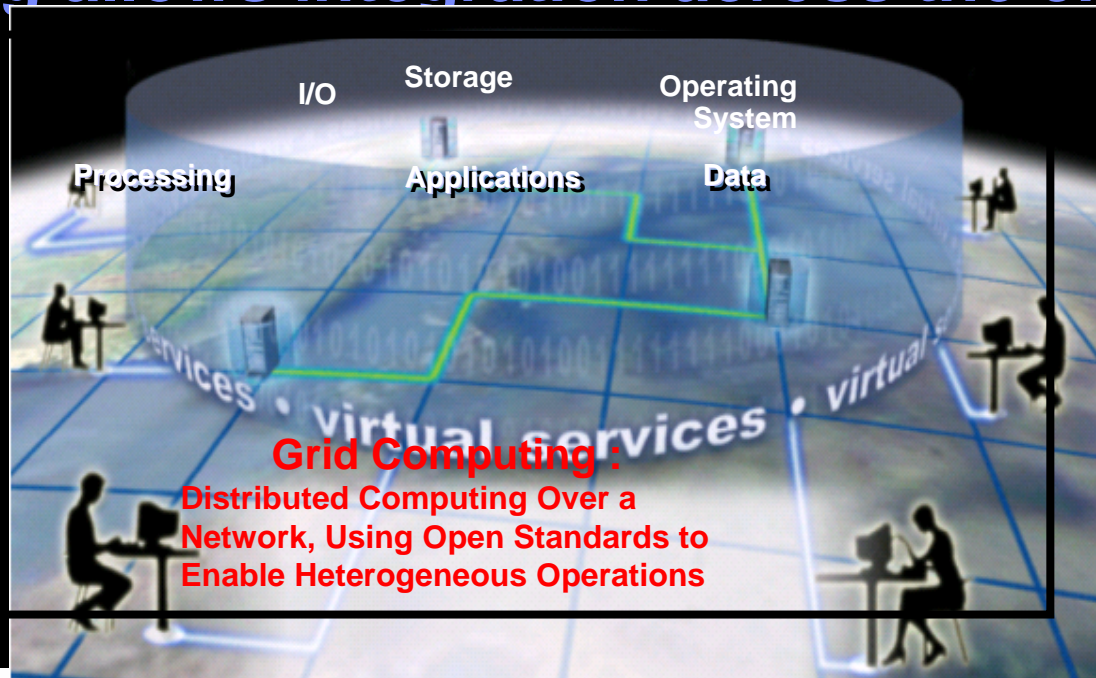
■"The Grid"

- ► **Enables management & sharing of and access to resources**
- ► **Delivers distributed computing over a network, using open standards to enable heterogeneous operations**
- ► **OGSA standards will be on all eServer systems**
- ► **Utilizes autonomic computing attributes to enhance manageability of Grids**
- ► **Leverages WebSphere cross-platform capabilities**

Grids make a safer world...

Disruptive technology to change client business model

Benefits:

Deutsche Bank

- Major reduction in processing time
  - 12 H to < 20 mn
- Significant reduction in TCO
- Intra-day calculations to reduce Market Risk Exposure
- Increased P&L opportunities

12 hours → 20 mn

Benefits:

IPK Institut für Pflanzengenetik und Kulturpflanzenforschung GATERSLEBEN

- Before Grid : 200 days for one single comparison of the 400,000 barley ESTs with the rice genome data
- With Grid : IPK can complete a one-time comparison in just 30 hours and a previously unfeasible project can now be completed in just a few weeks

200 days → 30 hours

Benefits:

charles SCHWAB

- **Processing time from four minutes to fifteen seconds**
- **Increase customer satisfaction**
- **Provide more robust wealth management applications**

4 mn → 15 s

Benefits:

MAGNA STEYR more value • more car

- Significant performance improvement (72 – 4 hrs)
- Risk and Error Reduction
- Cost Red. : Increased data accuracy Improved Time to Market

72 hours → 4 hours

Virtualization with Automation
Resource Sharing can help Increase Flexibility, Improve Service Levels And Reduce Costs

Virtualized Environment

Business Unit Executive — *Virtualization can help improve cost and speed*

Virtual Application Servers
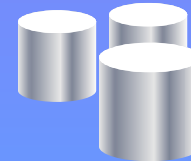
Virtual Storage

Virtual Networks

Systems Administrator — *More easily provision, manage and configure systems*

Virtualization

Physical Environment

Data Center Operations — *Make changes transparent to users*

Sun

HP

EMC

Hitachi

HP

Dell

Network Hardware

xSeries, BladeCenter, pSeries, iSeries, zSeries, Enterprise Storage Server, FAStT

IBM Virtualization Engine™ Integrated Packaging

Virtualization Systems Technologies

| Hypervisor |
| VLAN |
| Virtual I/O |

Integrated in Systems HW
Ordered as hw feature

IBM eServer and TotalStorage

- zSeries
- iSeries
- pSeries

- xSeries*
- ESS*
- BladeCenter*

Virtualization Systems Services

| IBM Grid Toolbox |
| TPM for eServer |
| Enterprise Workload Manager |
| Director Multiplatform |
| VE console |

Integrated Solution Building Block

IBM Virtualization Engine Suite for Servers

Support for AIX, i5/OS, z/OS*, agents

Support for Linux*, Windows & Solaris agents

Storage Virtualization Services

| Block Virtualization (SVC) |
| File Aggregation (SFS) |
| Storage Infrastructure Management (TPC) |

Integrated Solution Building Blocks
Ordered as separate products

IBM Virtualization Engine Suite for Storage

IBM TS SAN Volume Controller

IBM TS San File System
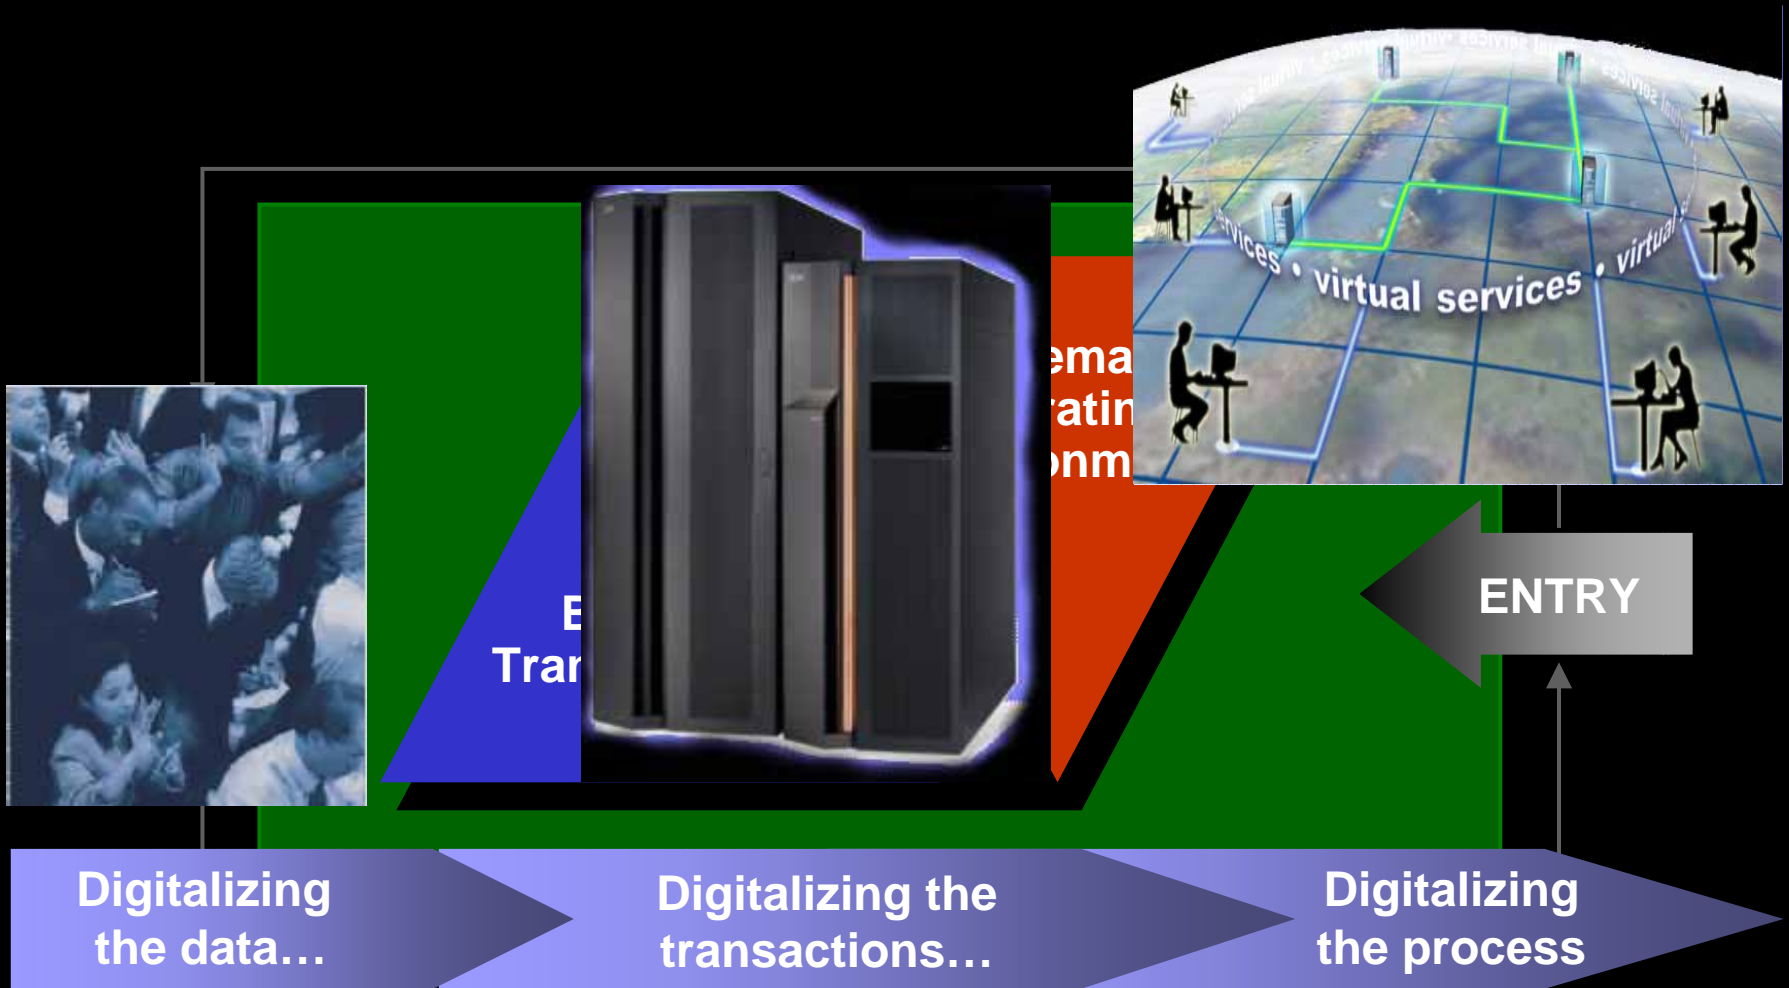
IBM TS Productivity Center

*not in first release

in short …

- **IBM leads in Core Technology**
- **IBM leads in Autonomic computing**
- **IBM leads in Grid Computing**
- **IBM leads in Linux**
- **IBM leads in Virtualization**
- **IBM leads in on demand**

To lead digitalization of processes in Pervasive Internet World

The Game : Digitalize the process

ema
ratin
onm

E
Tran

ENTRY

Digitalizing the data…

Digitalizing the transactions…

Digitalizing the process

IBM

Thank you

www.ibm.com/autonomic

IBM e-business on demand operating environment blueprint

Autonomic computing
Design self-managing systems and resources

IBM ecosystem for building self-managing systems

- **Technology**

- **Architecture**

- **Open standards**

- **Partners**

IGS Autonomic Readiness Offering

On demand operating environment

Business policy — The brain

Autonomic **core capabilities** — Spine and nerves

Products delivering autonomic features — The organs

Open | Integrated | Virtualized
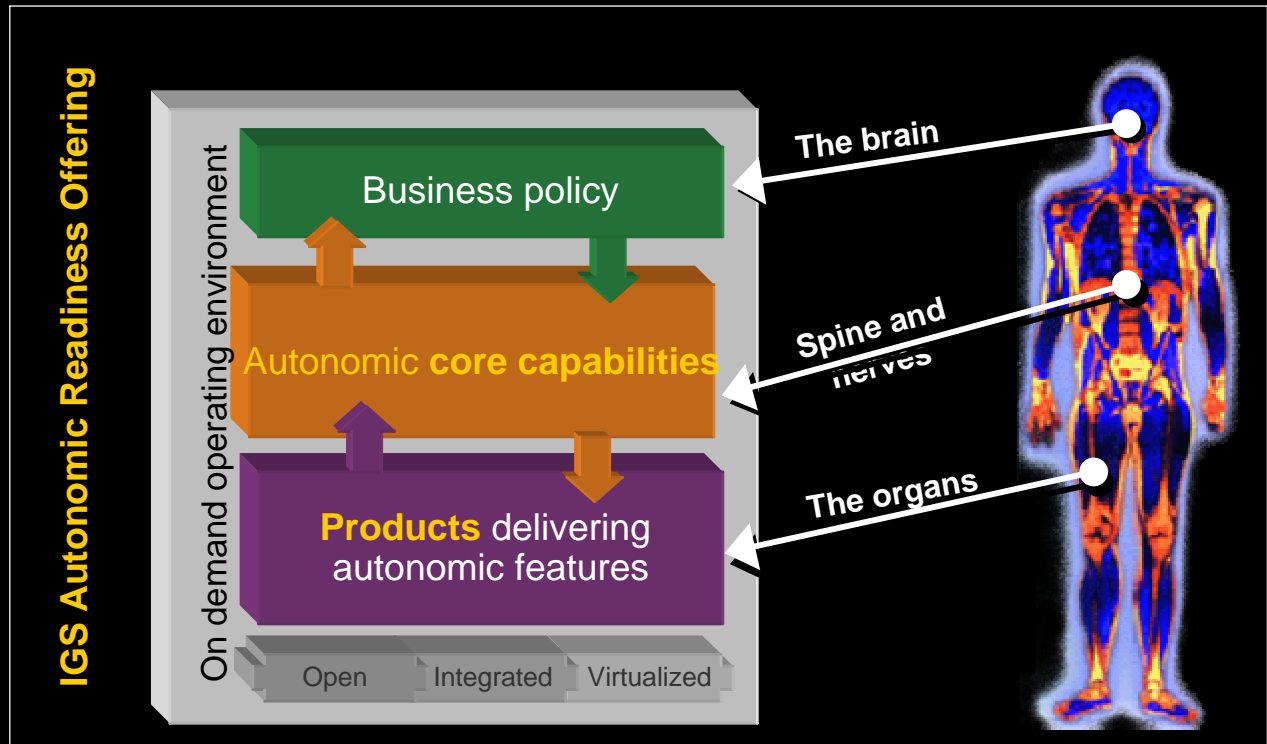
System Performance Stack

Performance improvements will increasingly require system level optimization

70%-80% CAGR, continuing

Virtualization, cluster management, etc., increasing

| Application |
| Middleware |
| S/W Development Tools | Operating System | Hypervisor |
| SMP System Structure
Fabric, switches, busses, memory system, protocols, … |

Uniprocessor approx. 50% CAGR, reducing

| Compilers |
| Microprocessor Core
Microacrchitecture, logic circuits, design methodology | Cache
Cache levels, granularity, latency, throughput | Interconnect |
| | | I/O |
| Package
I/Os, wiring level cooling, … |
| Semiconductors
Device, process, interconnect |

15%-20% CAGR, reducing

Application Server Provisioning Illustrated with zSeries and BladeCenter

Infrastructure Simplification

Orchestration

Linux
Linux
Linux   Linux
Linux   Linux   Linux
Linux

z/OS applications

Blade App svr
Blade DW app svr
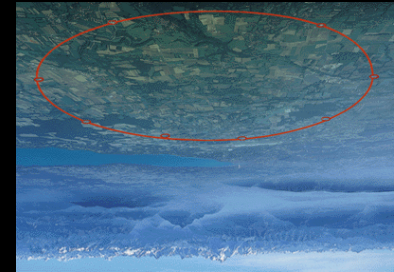Blade free-pool

z/OS Data Warehouse DB2
z/OS Data Server DB2
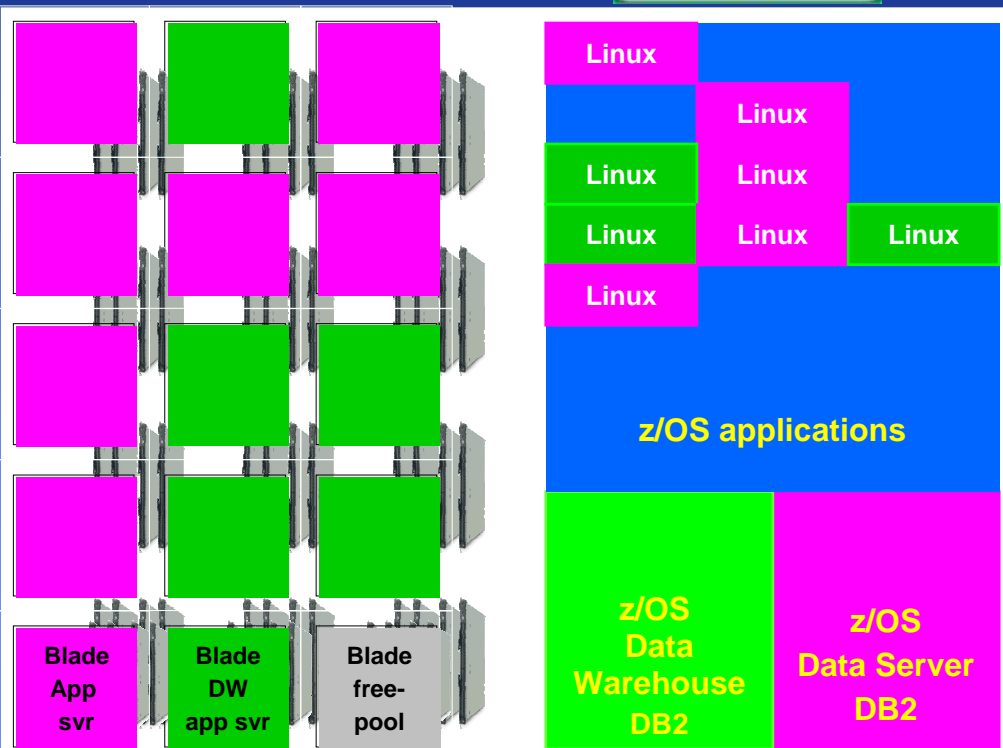
- SLA overachieved
- SLA overachieved
- Capacity reduction on blade pool and zSeries
- Provision reduction on blade pool
- Optimize zSeries resource utilization

Time1   Time2   Time3   Time4

SLA
SLA_{max}
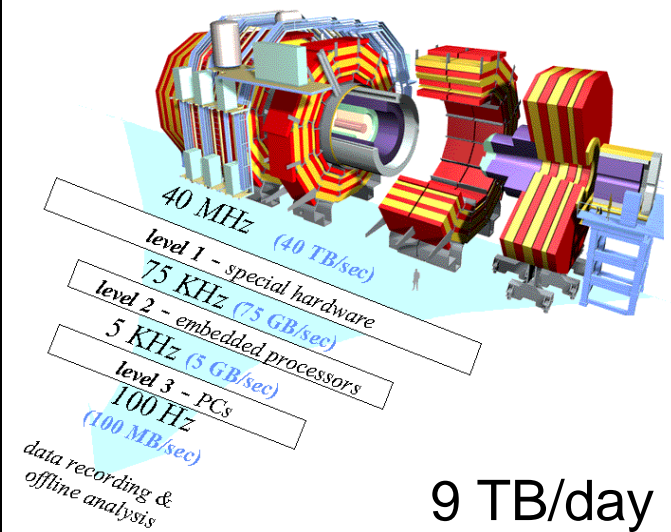SLA_{min}
time

Workload (WL)
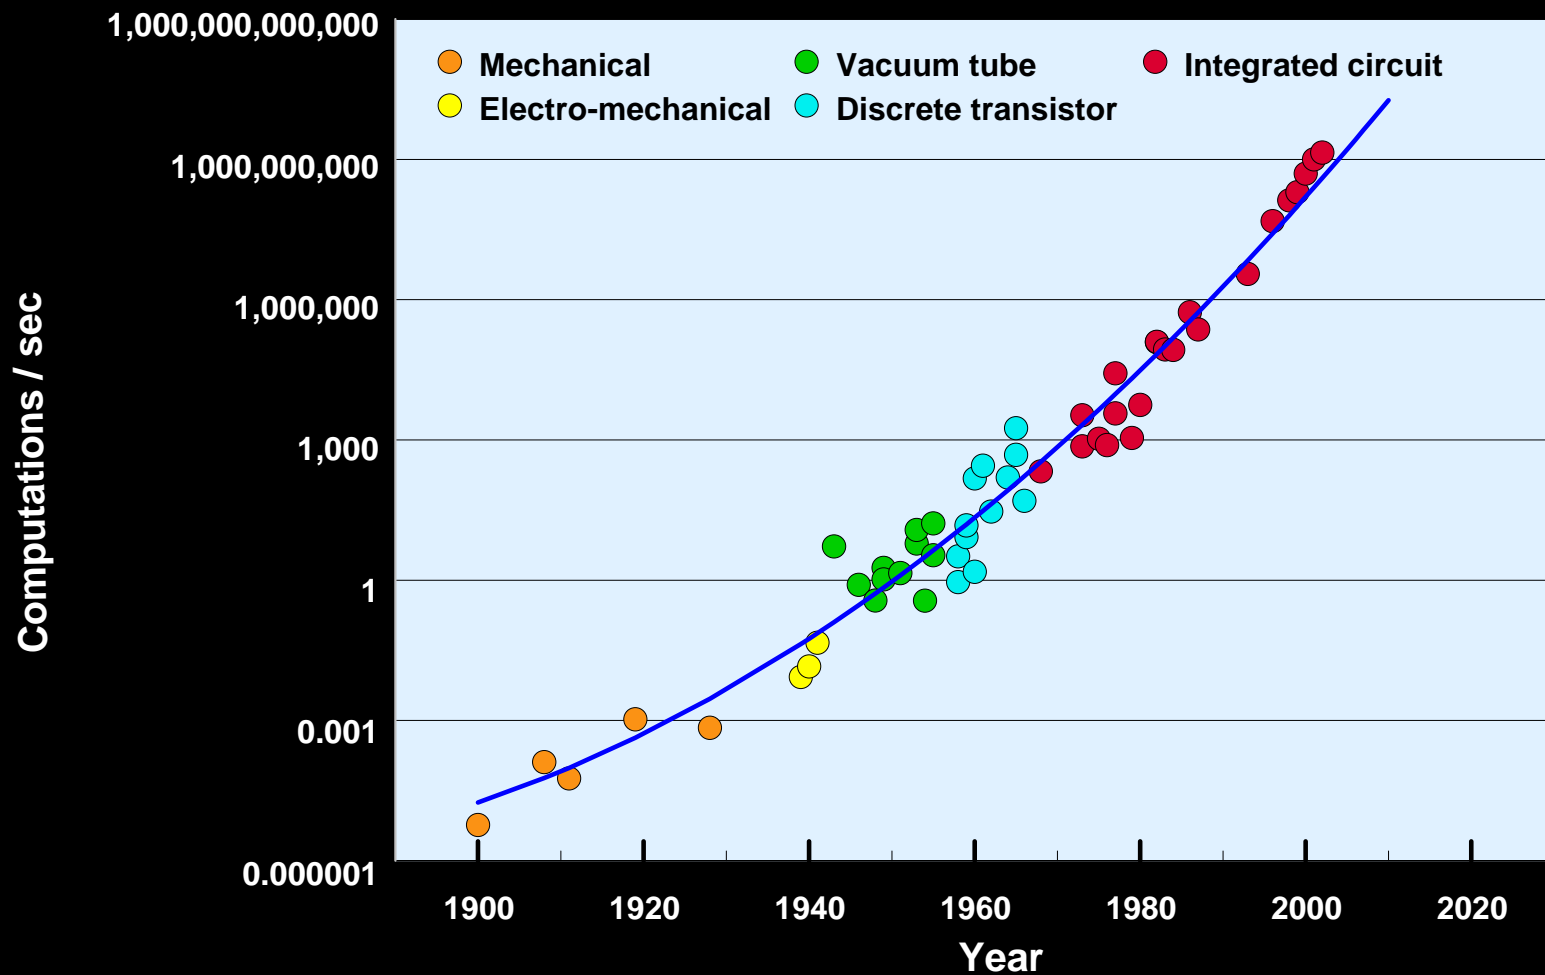blade/Cap_{max}
time
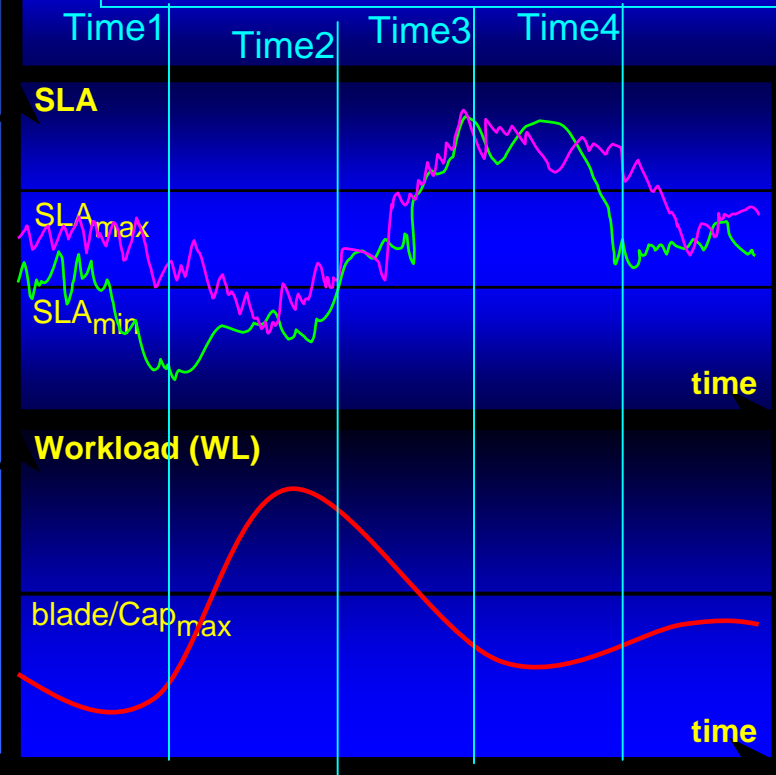
CERN - World's largest data Grid

- *April 02, 2003* : IBM - European Organization for Nuclear Research (CERN) collaboration

- Support of CERN's Large Hadron Collider (LHC) particle accelerator going live in 2007.

- World's largest data Grid

- IBM storage virtualization software -- Storage Tank

- To handle influx of 10 petabytes / a year by 2007

   the equivalent of 15 million CDs

- Only solution capable of handling both the capacity and the multi-platform requirements of the LHC.

- To recreate the initial moments our universe was formed - Understand nature of matter.

- Seamless access of more than 6000 scientists

Large Hadron Collider

40 MHz (40 TB/sec)
level 1 - special hardware
75 KHz (75 GB/sec)
level 2 - embedded processors
5 KHz (5 GB/sec)
level 3 - PCs
100 Hz (100 MB/sec)
data recording & offline analysis

9 TB/day

$1000 Buys

after Kurzweil, 1999 & Moravec, 1998