



Power within the **Grid**
Oracle **Grid** on **IBM** Servers

Authors:
Alain Roy
Paul Bramy
Oracle/IBM Joint Solutions Center

November 14, 2006



Agenda

- IBM/ORACLE & JSC.
- Shared Technical foundations.
- Customer Projects presentation.
- Conclusions.



EMEA Oracle/IBM
Joint Solutions Center
MONTPELLIER - FRANCE



The New IBM & Oracle Relationship

- IBM and Oracle have had a relationship for 19+ years (27 years with JDE, 16 with PeopleSoft, 8 with Siebel) >15,000 joint customers WW
- PSE&JDE are running Websphere & DB2.
- Renewed S&D relationship
- Enhanced strong BCS relationship -> N°1 Oracle Integrator
- IBM viewed by Oracle as a significant partner
- Strong technology relationship Development, Benchmarking & Performance: 6 joint solution centers worldwide
- Oracle Global Sales Support Team
- IBM participation in Oracle customer events
 - Open World, Oracle Regional User Group Events
- Oracle Participation in IBM Events like Unix in Focus, European System p Technical University
- 9i, 10g and 11i, all available on AIX 5.3 at GA
- Daily builds on AIX during development cycle (30 engineers at Redwood shore)
- Oracle Announcement to join the IBM AIX Collaboration Center (IACC) as a founding partner (*)



IBM

&

ORACLE

ORACLE

Oracle/IBM Joint Solutions Center

What is the JSC?

The Oracle/IBM Joint Solutions Center, created in 1998, is a presales team, permanently staffed by Oracle & IBM, to contribute to the successful delivery of Oracle/IBM Joint Solutions.

Why the JSC is unique?

- Located inside the PSSC, the IBM's European Benchmarking and Support center (PSSC) in Montpellier, France. So JSC has access to large amount of hardware from all the IBM brands (p, x, l, z).
- It is staffed by Subject Matter Experts from both IBM & Oracle so we can have access to the best of our two corporations simultaneously for the benefit of our joint customers.

Pre Sales Assistance through the whole project/sales cycle (covering more than 100 RAC Projects last Fiscal year)

- Help me to Convince: Briefings & Conferences (40/year)
- Help me to Build: Architecture, Design, Sizing Assistance (300/year)
- Help me to Demonstrate: PoC, Benchmarks Assistance (30/year)
- Help me to Deliver: hands-on workshops (5/year), Publications (10/year)

Infrastructure & Investments

People

- A unique IBM investment in Europe
- Skill set: 230 IT Specialists & Architects

Onsite customer facilities

- 4 machine rooms + 1 TotalStorage lab
- 7 briefing rooms (12 to 150 seats)
- 11 education rooms
- 18 benchmarking rooms

Remote customer facilities

- 2 large platforms (beta test & FMS)
- 6 midsize platforms
- 2 large HPC Clusters
- Access to the GDPS solution center platform

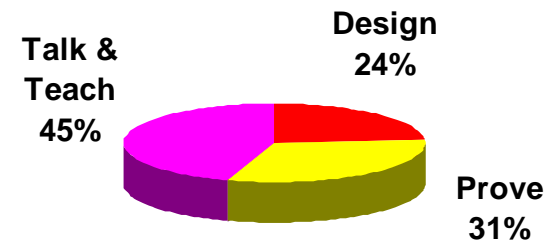
Latest Technologies

IBM System z	197,000 Mips	284 Terabytes
IBM System p	18 Teraflops	92 Terabytes
IBM System i	85,000 CPW, 15 servers	40 Terabytes
IBM System x & Blades	408 servers	Intel EM64T, AMD Opteron, BladeCenter-E/T, x366, x460
IBM System Storage	40 ESS/DS8000	DS4000/DS6000 & 1300 fibre channel ports, including FICON


Activities & Clients in 2006



1046 Activities



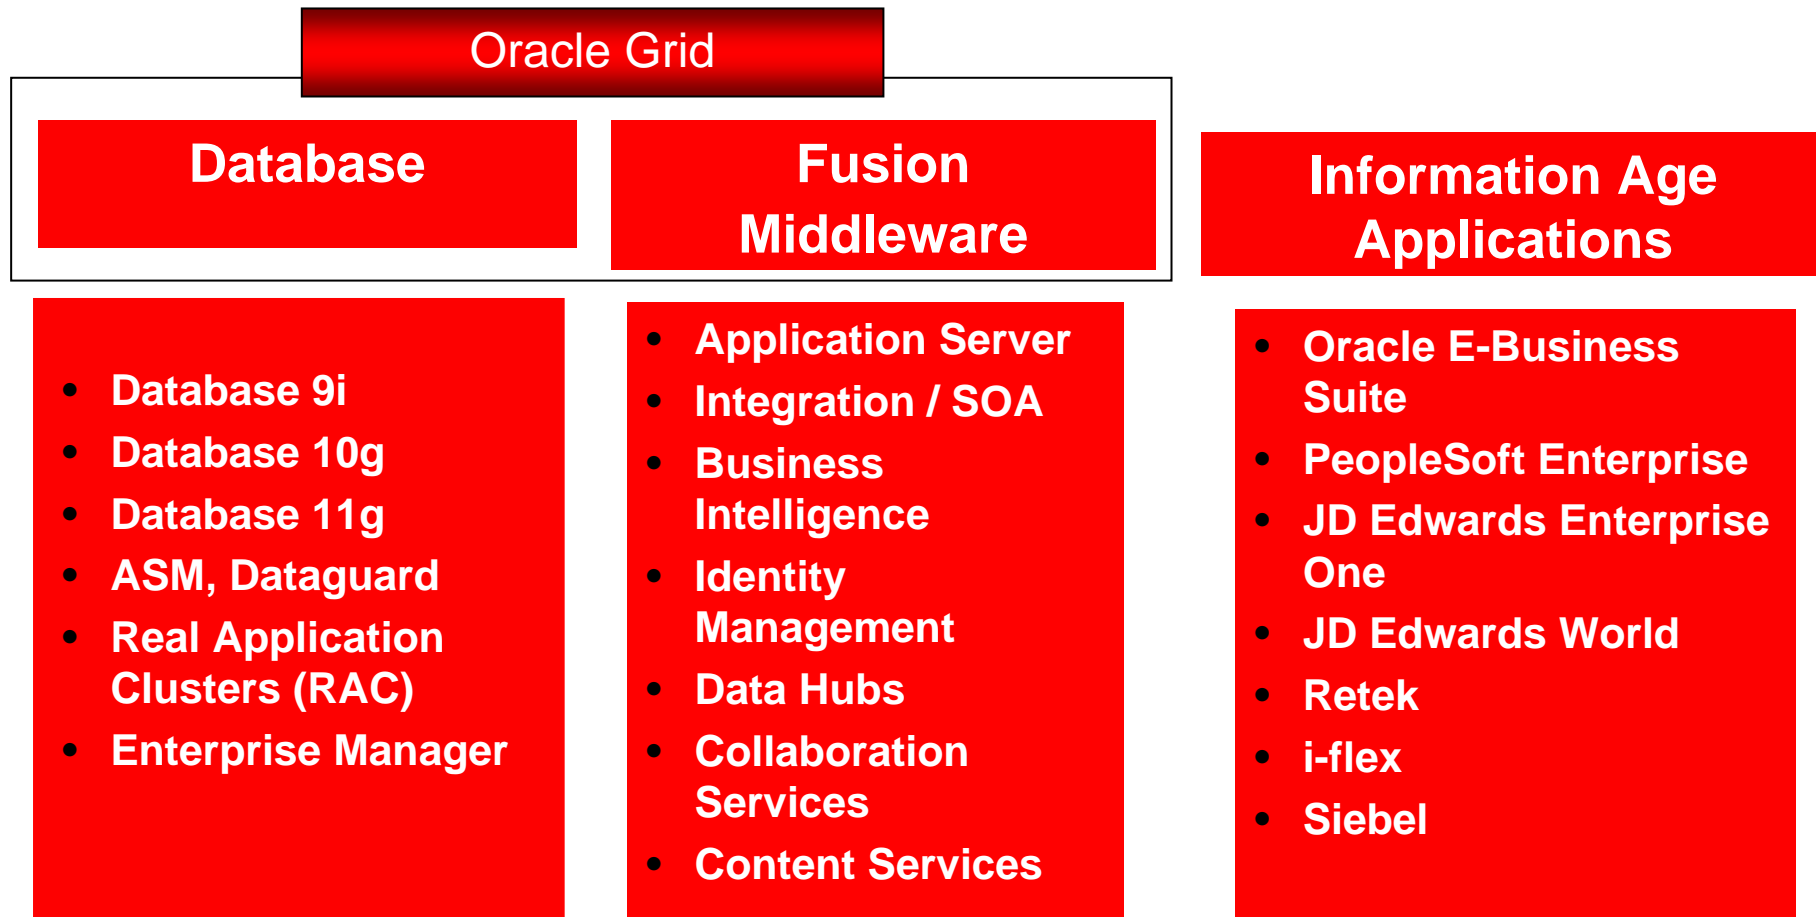
7142 Visitors from 1190 Companies



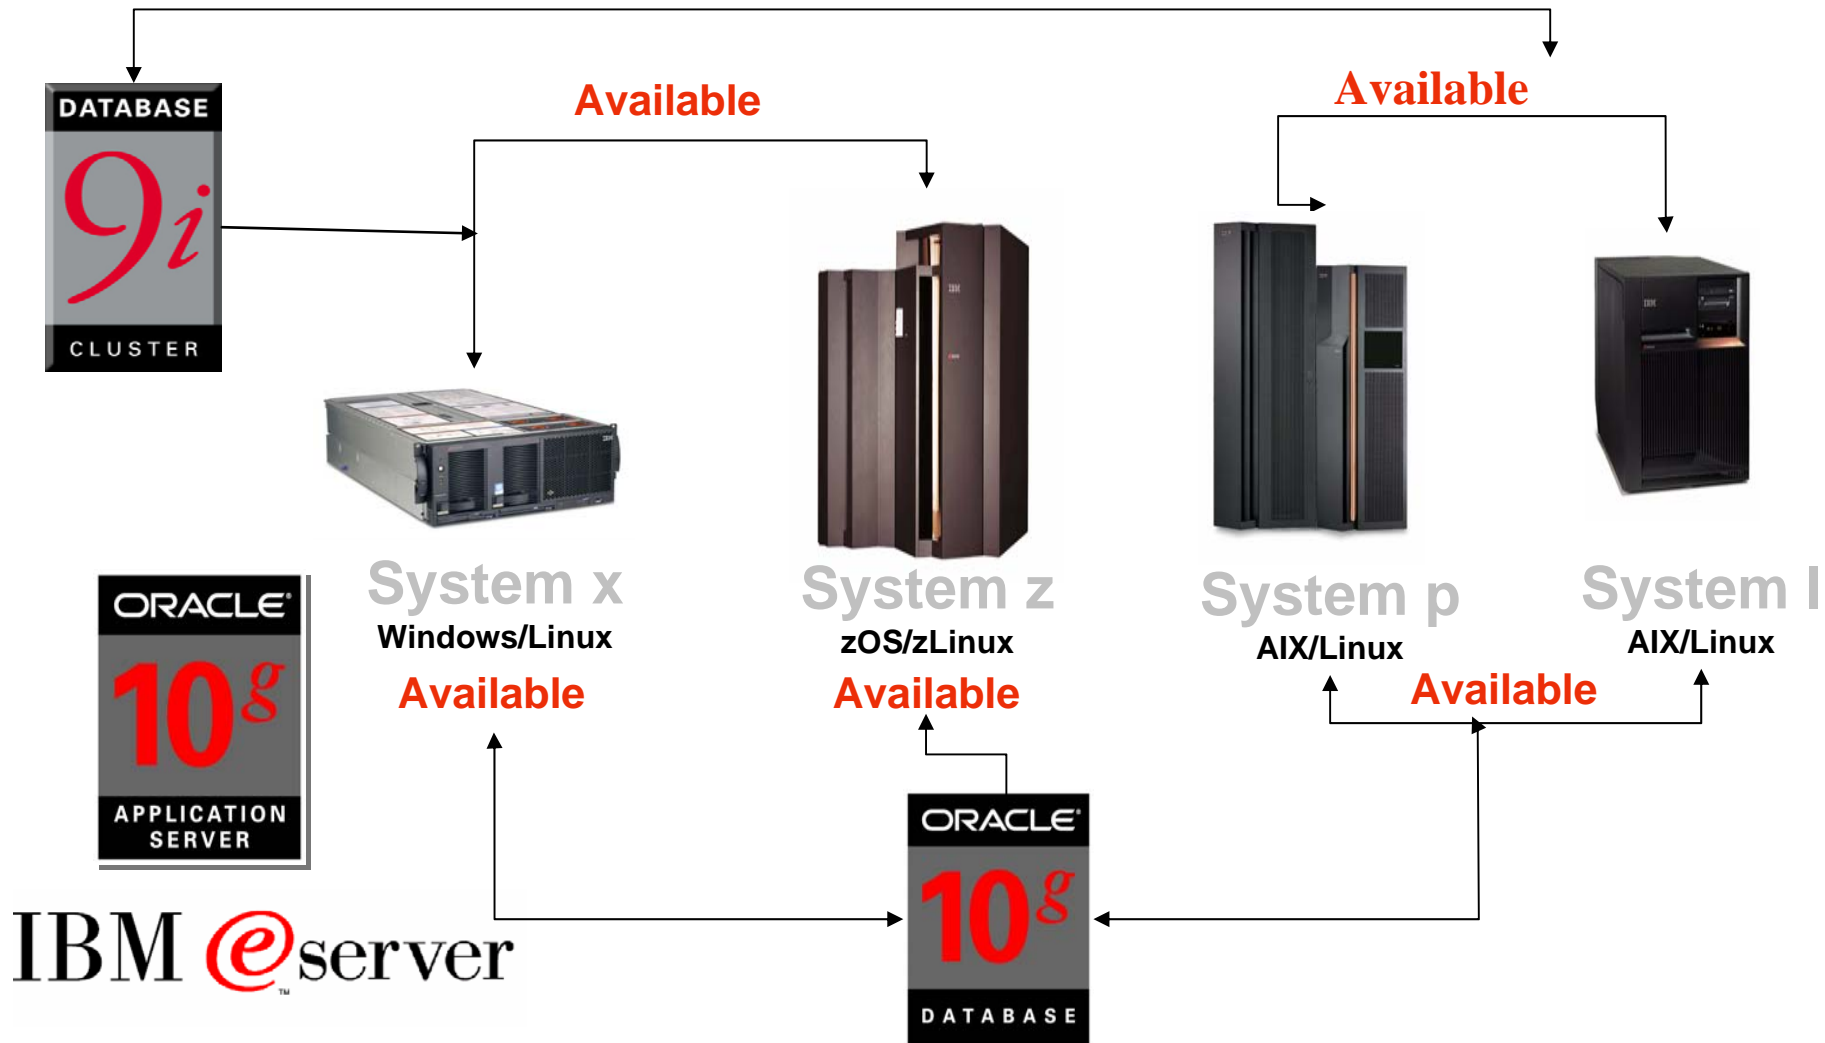
JSC & Grid current deliveries

- **One-to-Many or One-to-One Workshops**
 - Implementing RAC on System z running Linux, on System p running AIX5L or Linux, on System x running Linux
 - Implementing Oracle EBS 11i RAC on System p running AIX5L
 - **New in Dec 06 : Implementing Oracle 10g R2 RAC for SAP on System p running AIX5L !**
 - **New in Feb 2007 : Same but on Linux on Power !**
- **Best Practices Cookbooks**
 - Implementing Oracle 10g R2 RAC on System p running AIX5L, System x running Linux, System z running Linux & System p running Linux.
- **Customer Benchmarks : Qualification & Execution**
- **Certification**
 - OVC for Oracle 10g RAC with ASM Mirroring on System x running Linux/.
 - NFS Certification for 3rd Voting Disk with AIX5L Nodes
 - Considered: Oracle 10g RAC & ASM on System p with Virtualized IOs (VIOs)
 - OASB Sizing Benchmark on EBS / System z / Linux
- **Share JSC Infrastructure for your practicing**
- **Awareness Events & Conferences**
 - Oracle/IBM Roadshows (Denmark, Russia, Romania, Middle East, South Africa and next week Italy)
 - Large Events such as Unix in Focus, CSI Interchange – Montpellier, Open World (USA),....
 - One to One briefings in Montpellier or onsite
- **Design Workshops**

Oracle products overview



Oracle Technology from the Blade to the mainframe...



Oracle Database on IBM eServers



- Oracle 6 (1989) Already available on AIX, MVS
- Oracle 7 (1994) AIX: Leader in OPS/Clusters performance Available on AIX, MVS, VM
- Oracle 8.0 (1998) AIX: "Post-wait" system call & AIX: "kstat" performance monitoring Available on AIX, OS390, Windows
- Oracle8i (2000) Improved Async I/O/ AIX: Support for Hardware multi-threading / Full HACMP/ES CRM integration Available on AIX, OS390, Windows, intel (Terminal version 8.1.7)
- Oracle9i release 1 (2001) AIX: Full RAC support; RAC Guard support; Available on AIX, zOS, Windows, Linux/xSeries Current version 9.0.1.
- Oracle9i release 2 (2002) : AIX: Exploits new Power4 features, Large page support, 64-bit kernel; Introduction of LPAR, True hardware partitioning, JFS2 ... petabyte-size FS Available on AIX 4.3.3, and AIX 5L (5.1,5.2,5.3) zLinux / zOS, Windows, Linux x86

A history of working together

Oracle10g a Joint effort



*Already
working on
Oracle 11g !!!*

A 20-month project (June 2002 to February 2004)

- Concurrent development on AIX 5L
- Daily build on AIX 5L concurrently with the base platforms
- Additional projects in parallel of regular builds (TPC-C, ...)

To Jointly Improve Oracle10g/AIX5L

- Codepath analysis on POWER4 – feedback to Oracle Server Technologies
- Stress-testing with artificial and customer workloads
- Reducing per-process private memory footprint
- VisualAge v6
- Direct line to IBM Austin performance teams

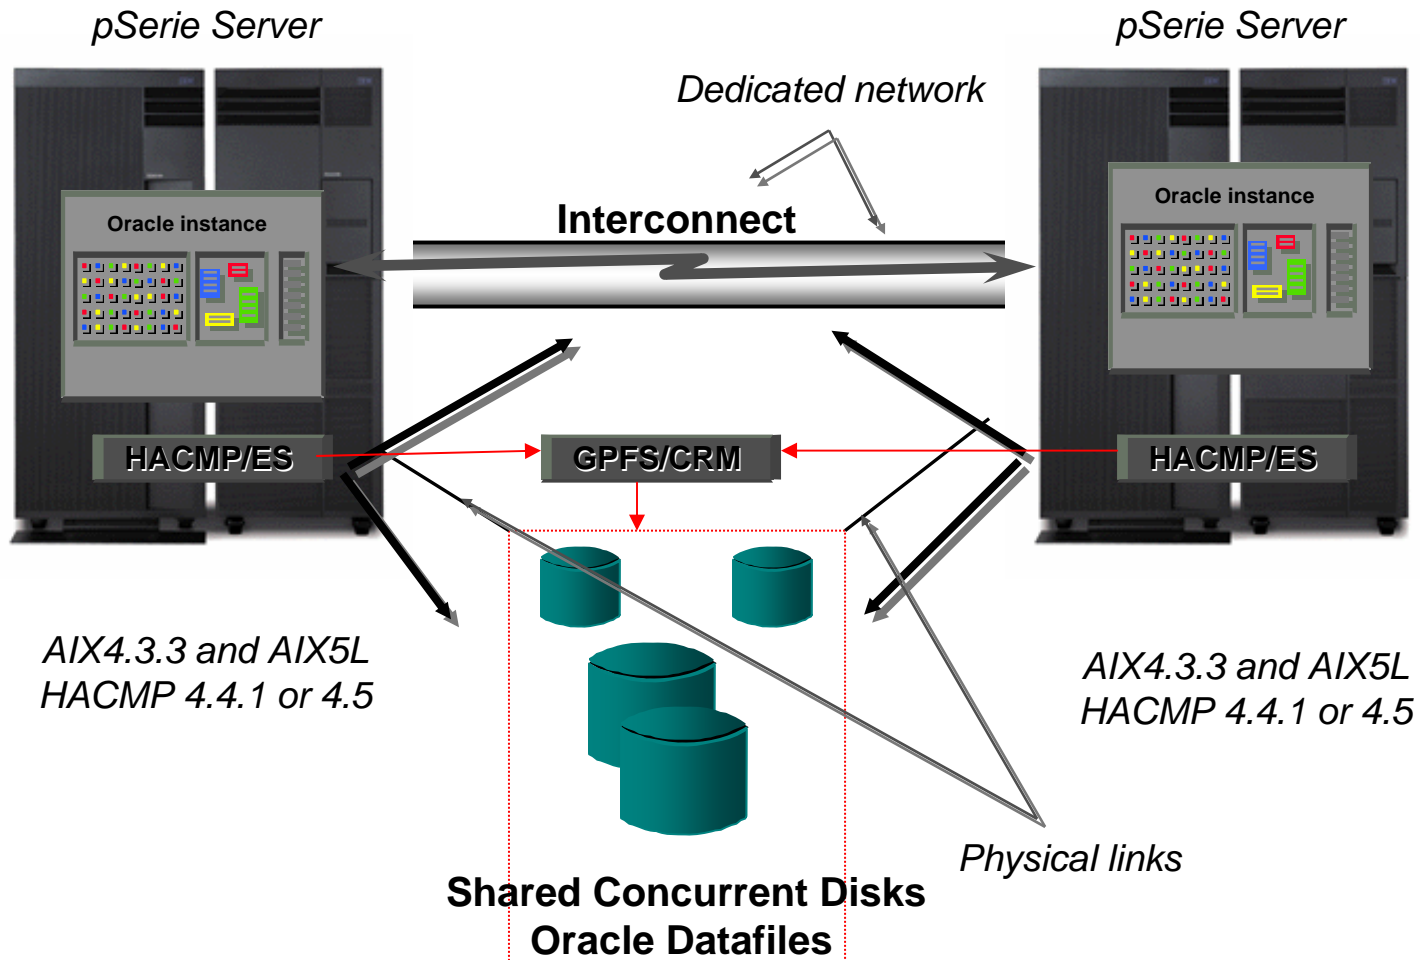
With an emphasis on RAC

- New Cluster Topologies & RAC interconnect changes
- Integration of GPFS to Oracle development and test cycle

Main Goals

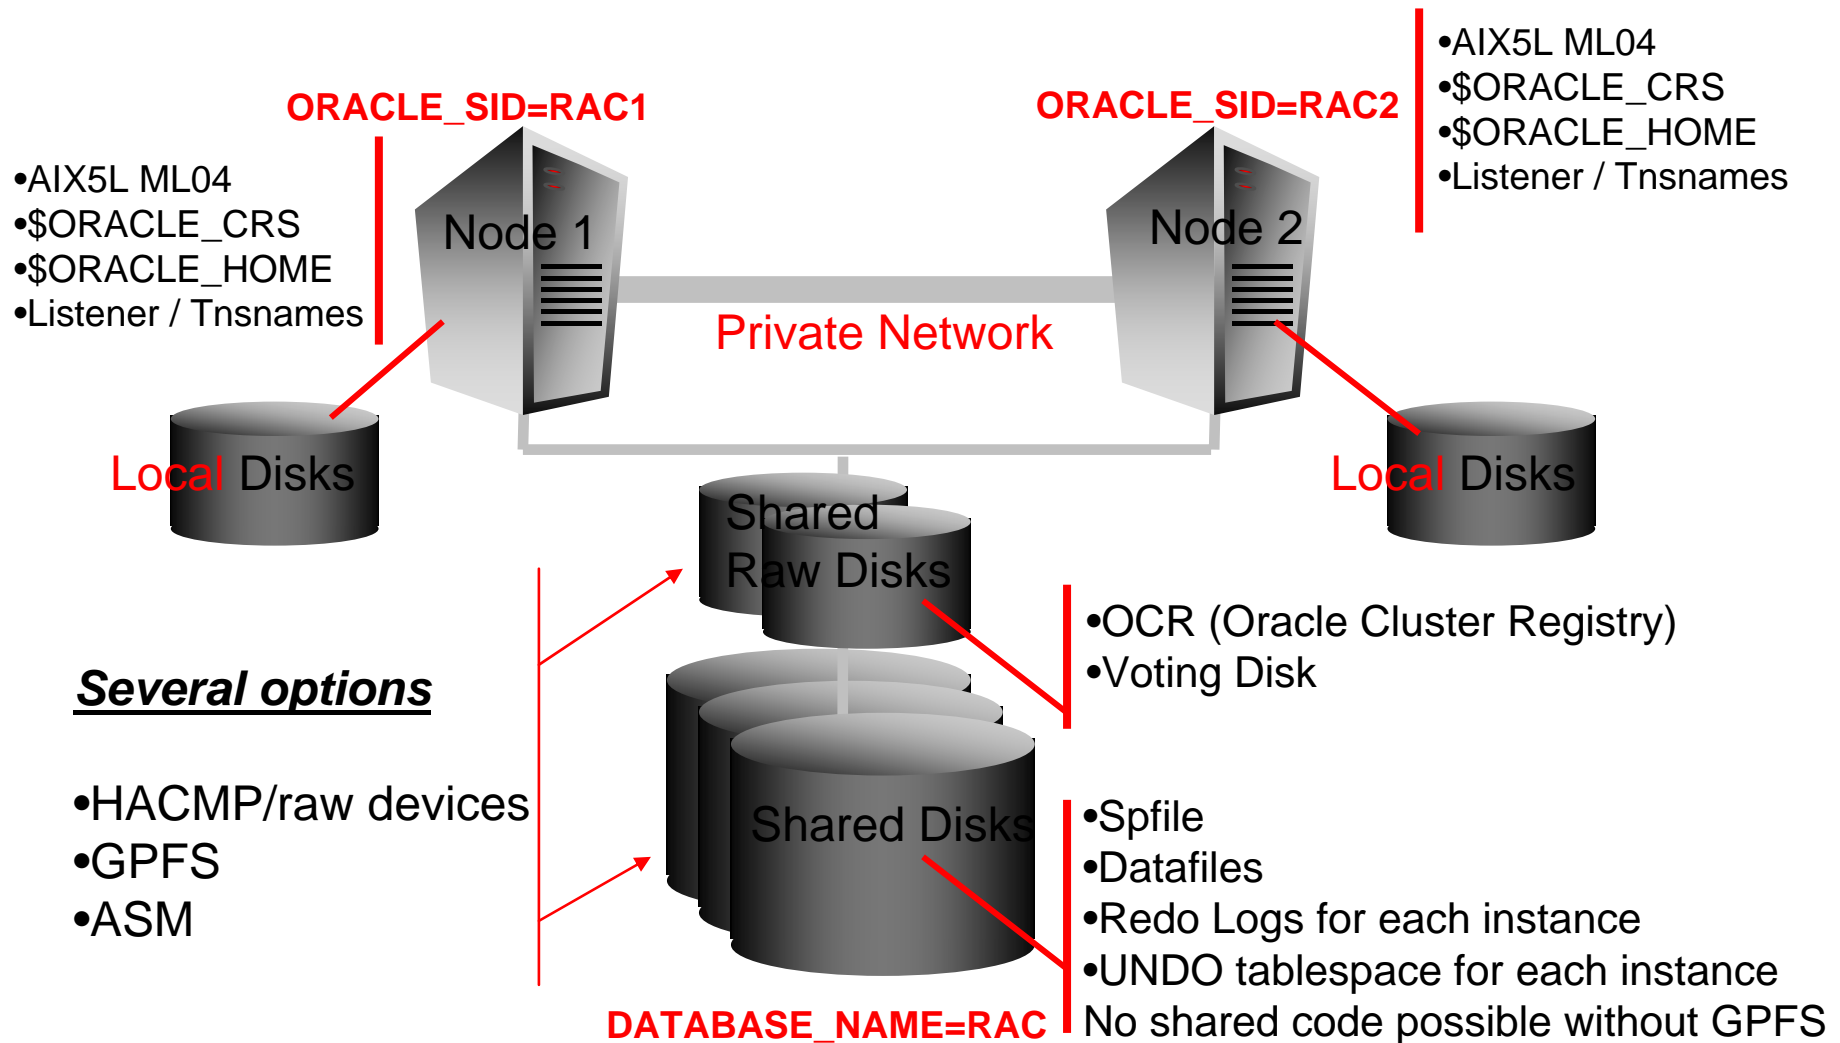
- Make sure Oracle uses the best and latest pSeries & AIX features
- Simplify File Systems performance, I/O performance tuning
- Simplify RAC configuration and install

Oracle9i RAC Architecture on IBM pSeries

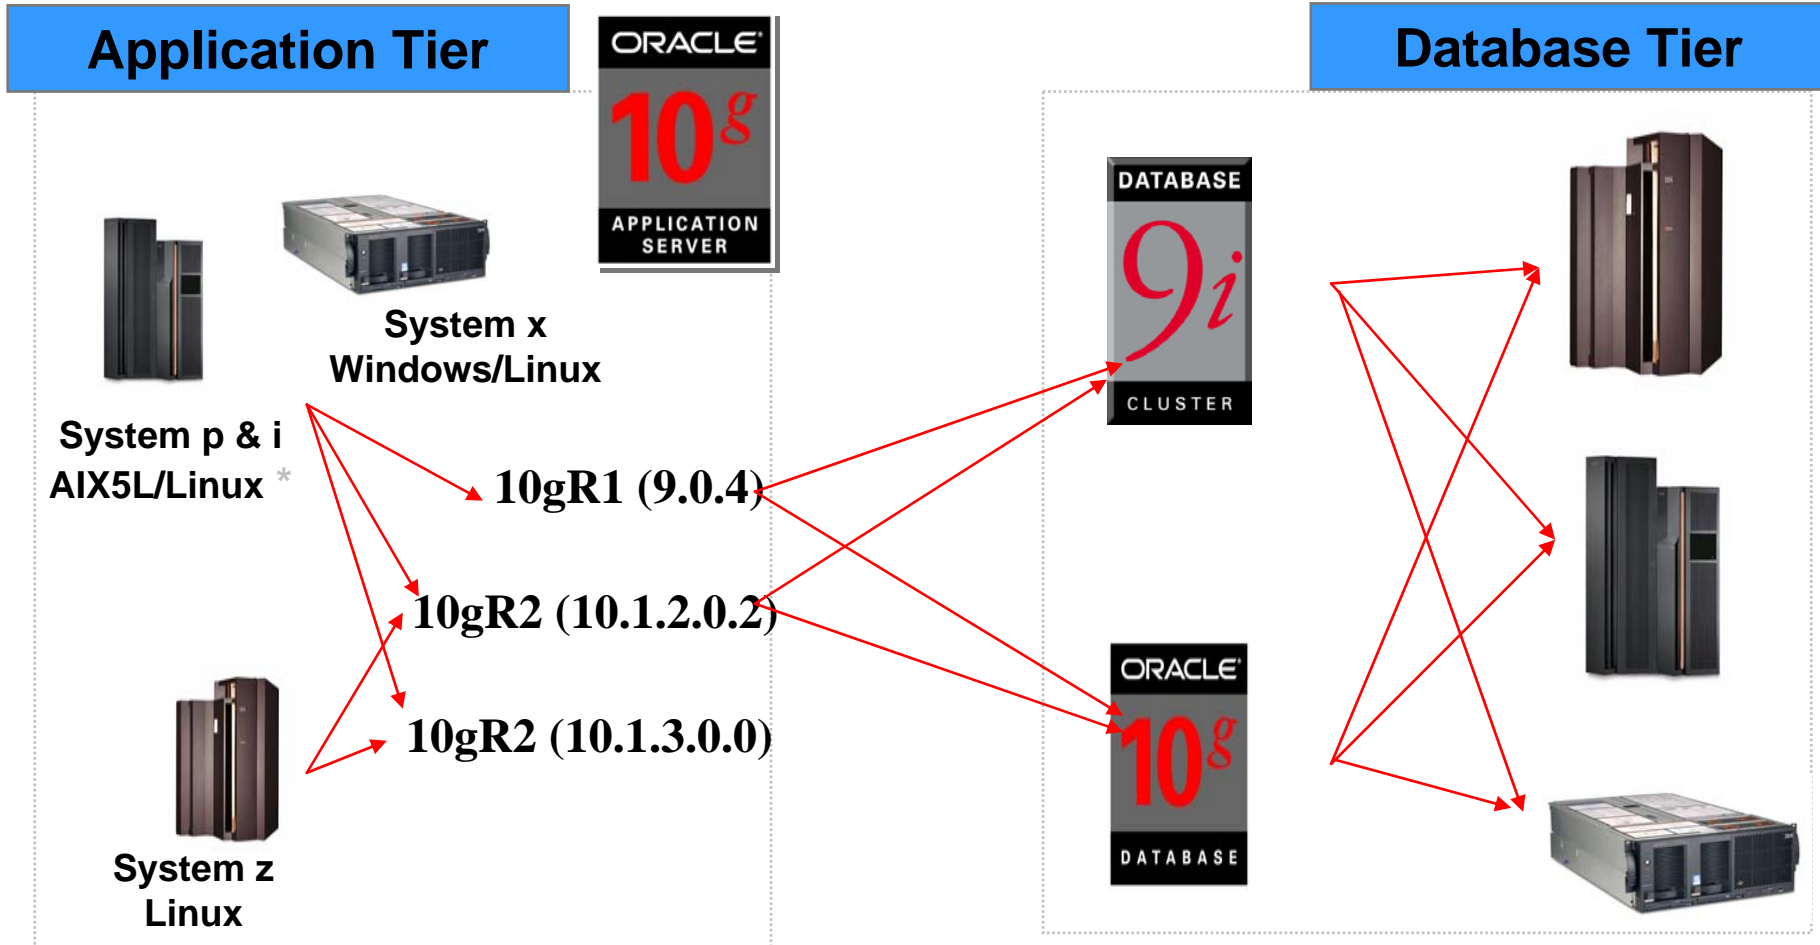


IBM Storage sub-systems and other Disks vendors

Oracle 10g RAC architecture on AIX5L



Oracle Application Server on IBM eServers



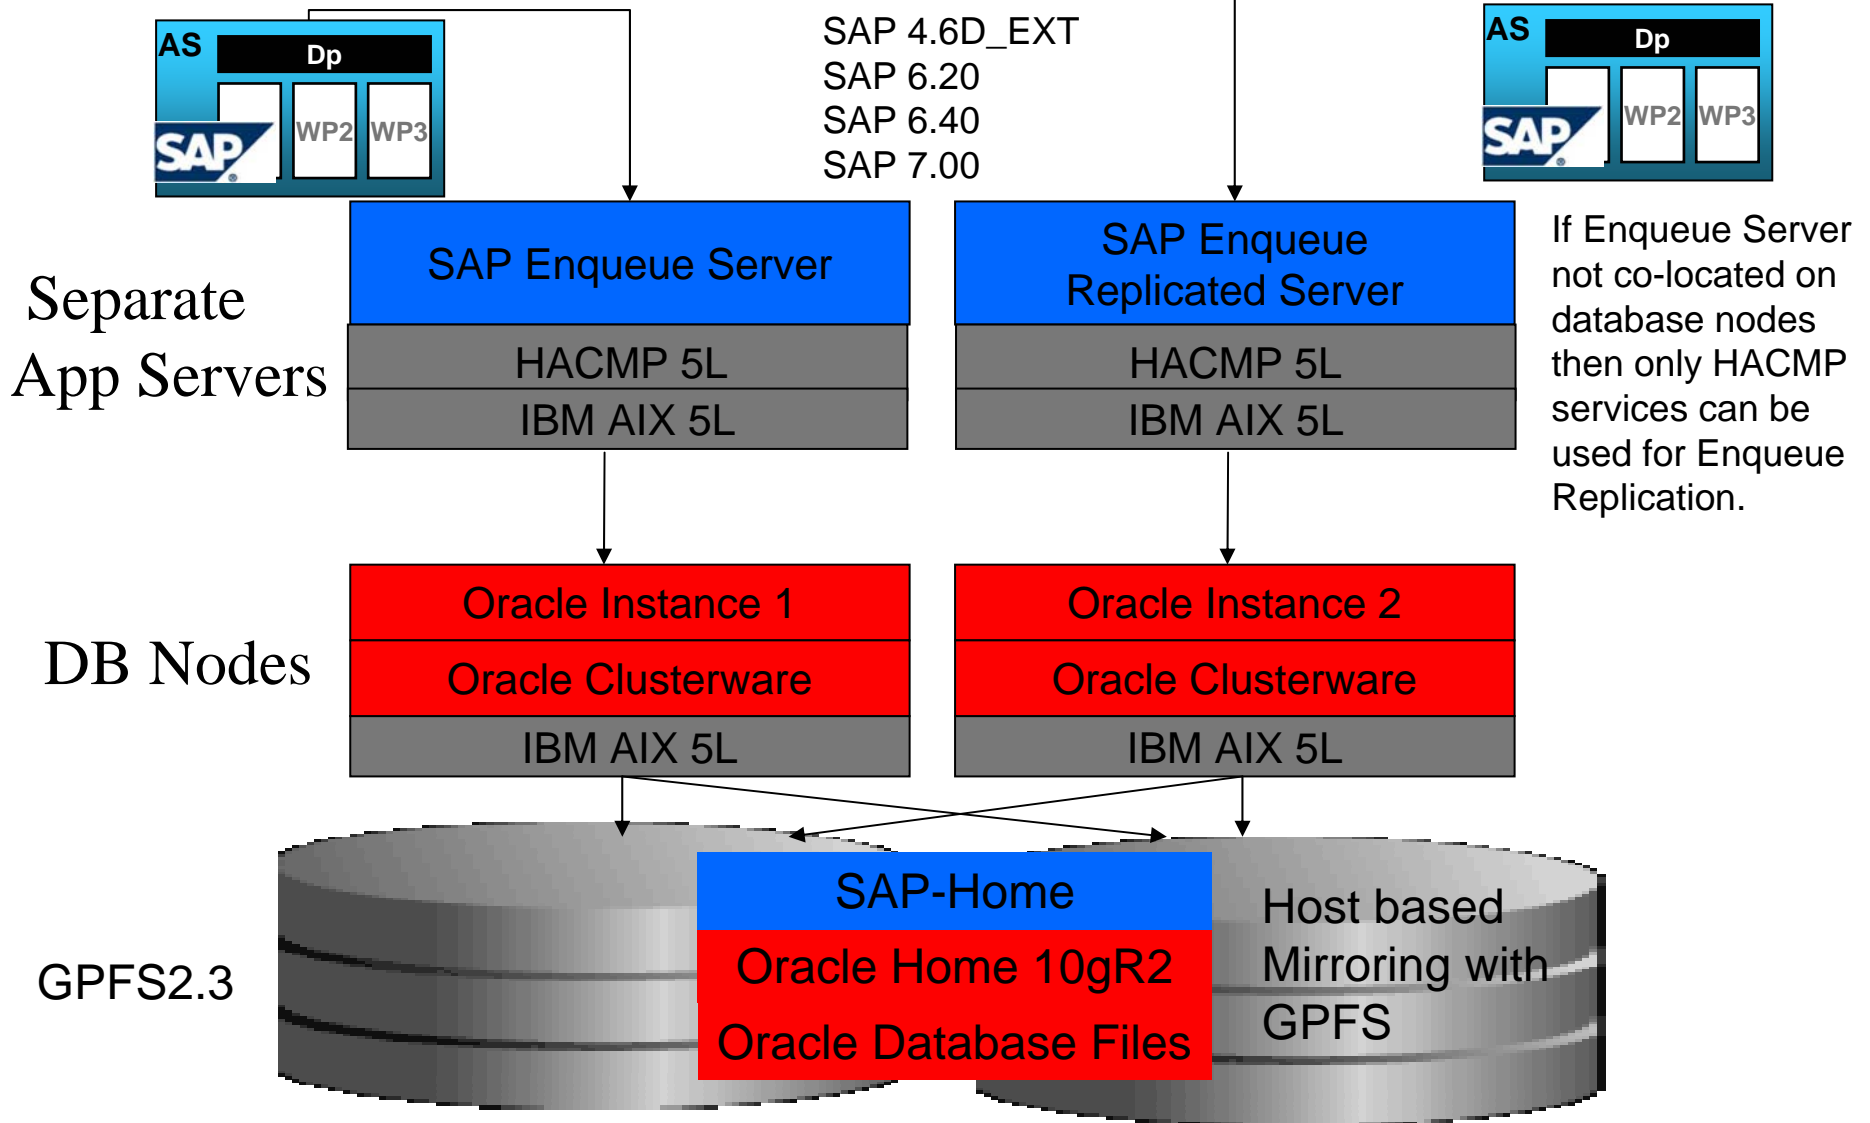
- Oracle 10g AS is certified with Peopletools 8.47

All major Oracle Applications run on IBM eServers

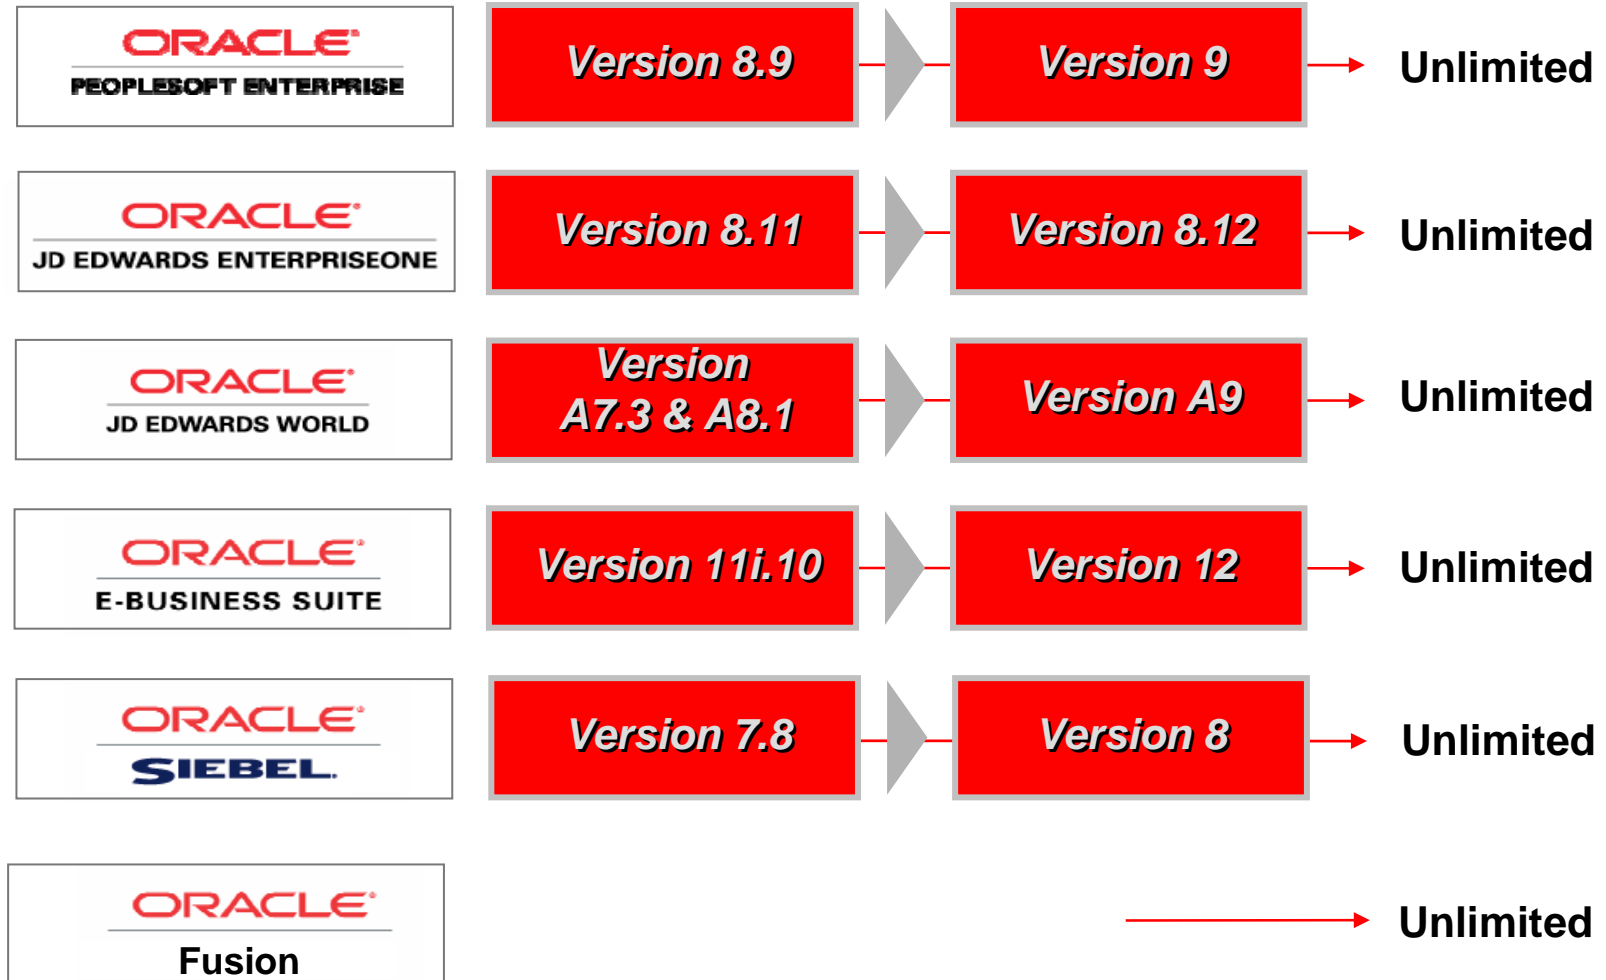


Oracle E-Business Suite		(AIX partition)
Oracle OCS		(AIX partition)
Oracle Technology: Oracle Database & Application Server		(AIX partition)
PeopleSoft Enterprise		(AIX partition)
JD Edwards EnterpriseOne		
JD Edwards World		
Siebel		(AIX partition)
Retek		(AIX partition)
ProfitLogic		(AIX partition)
i-flex		(AIX partition)
360Commerce		(AIX partition)

SAP Solutions on IBM AIX using Oracle RAC 10gR2

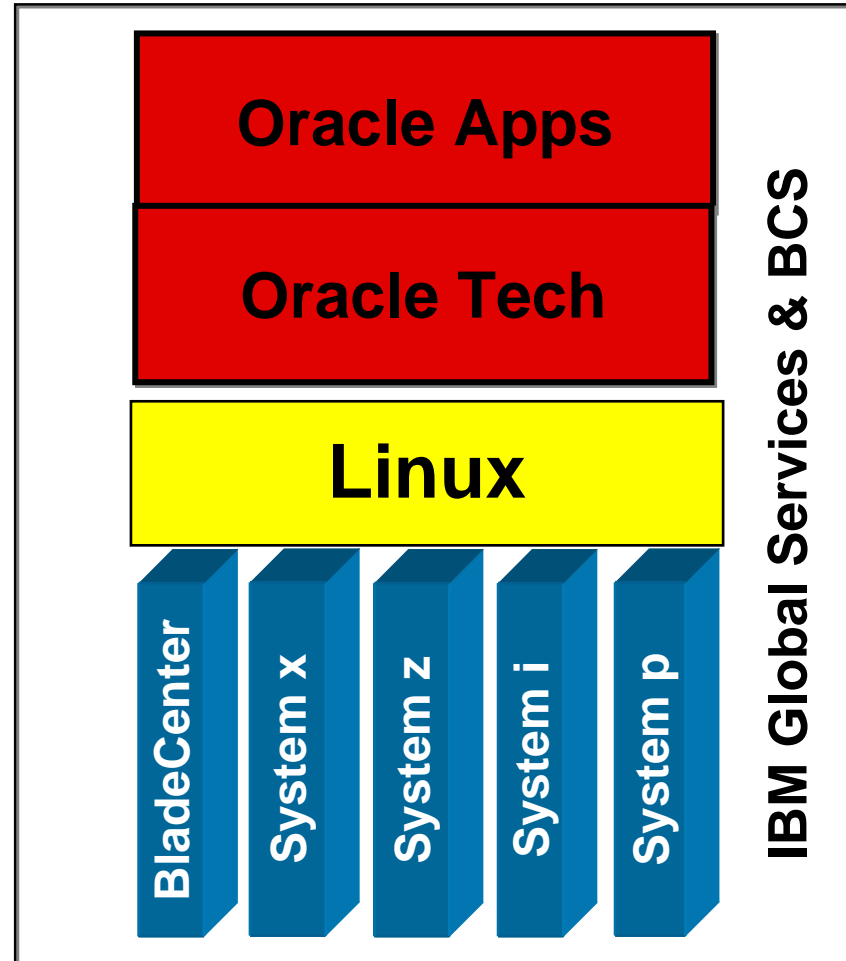


Oracle Application Roadmap




An other good reason of the teaming

Linux Now Runs On
All
 IBM eServer
 Platforms



ORACLE®

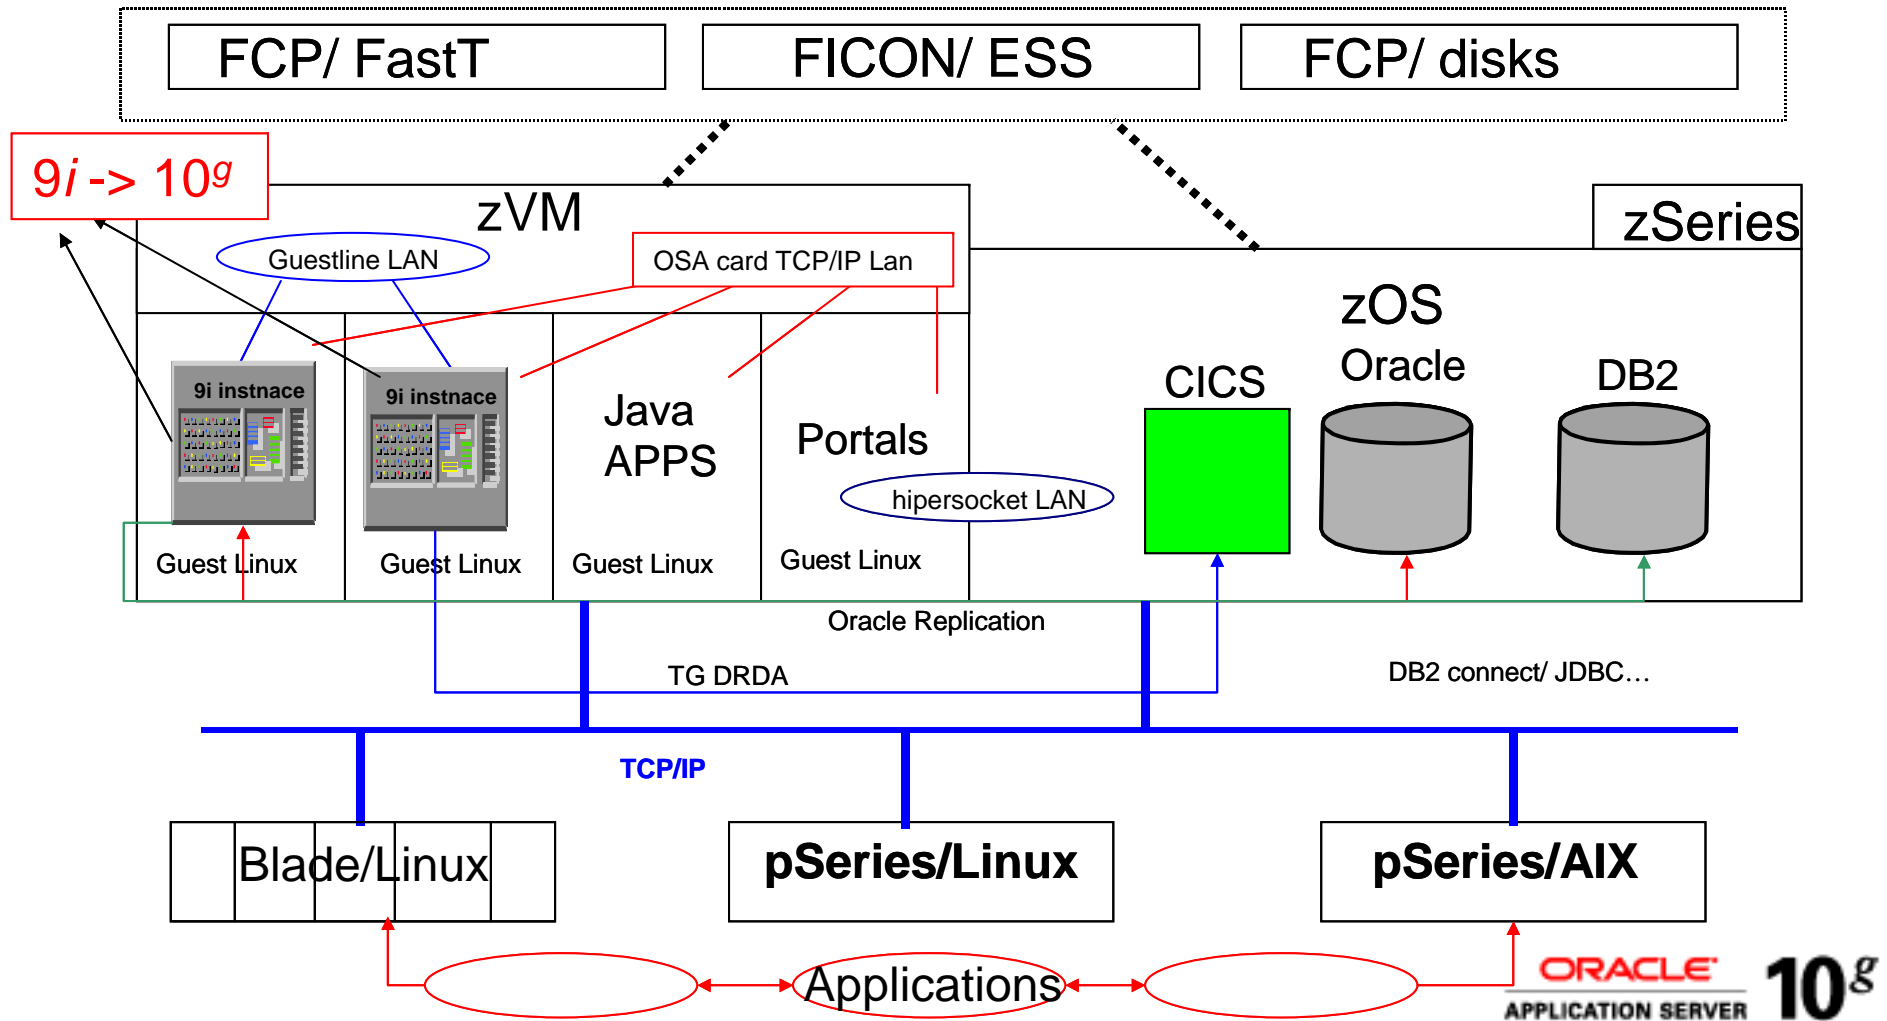


IBM and Oracle collaborate to grow mainframe Linux workload

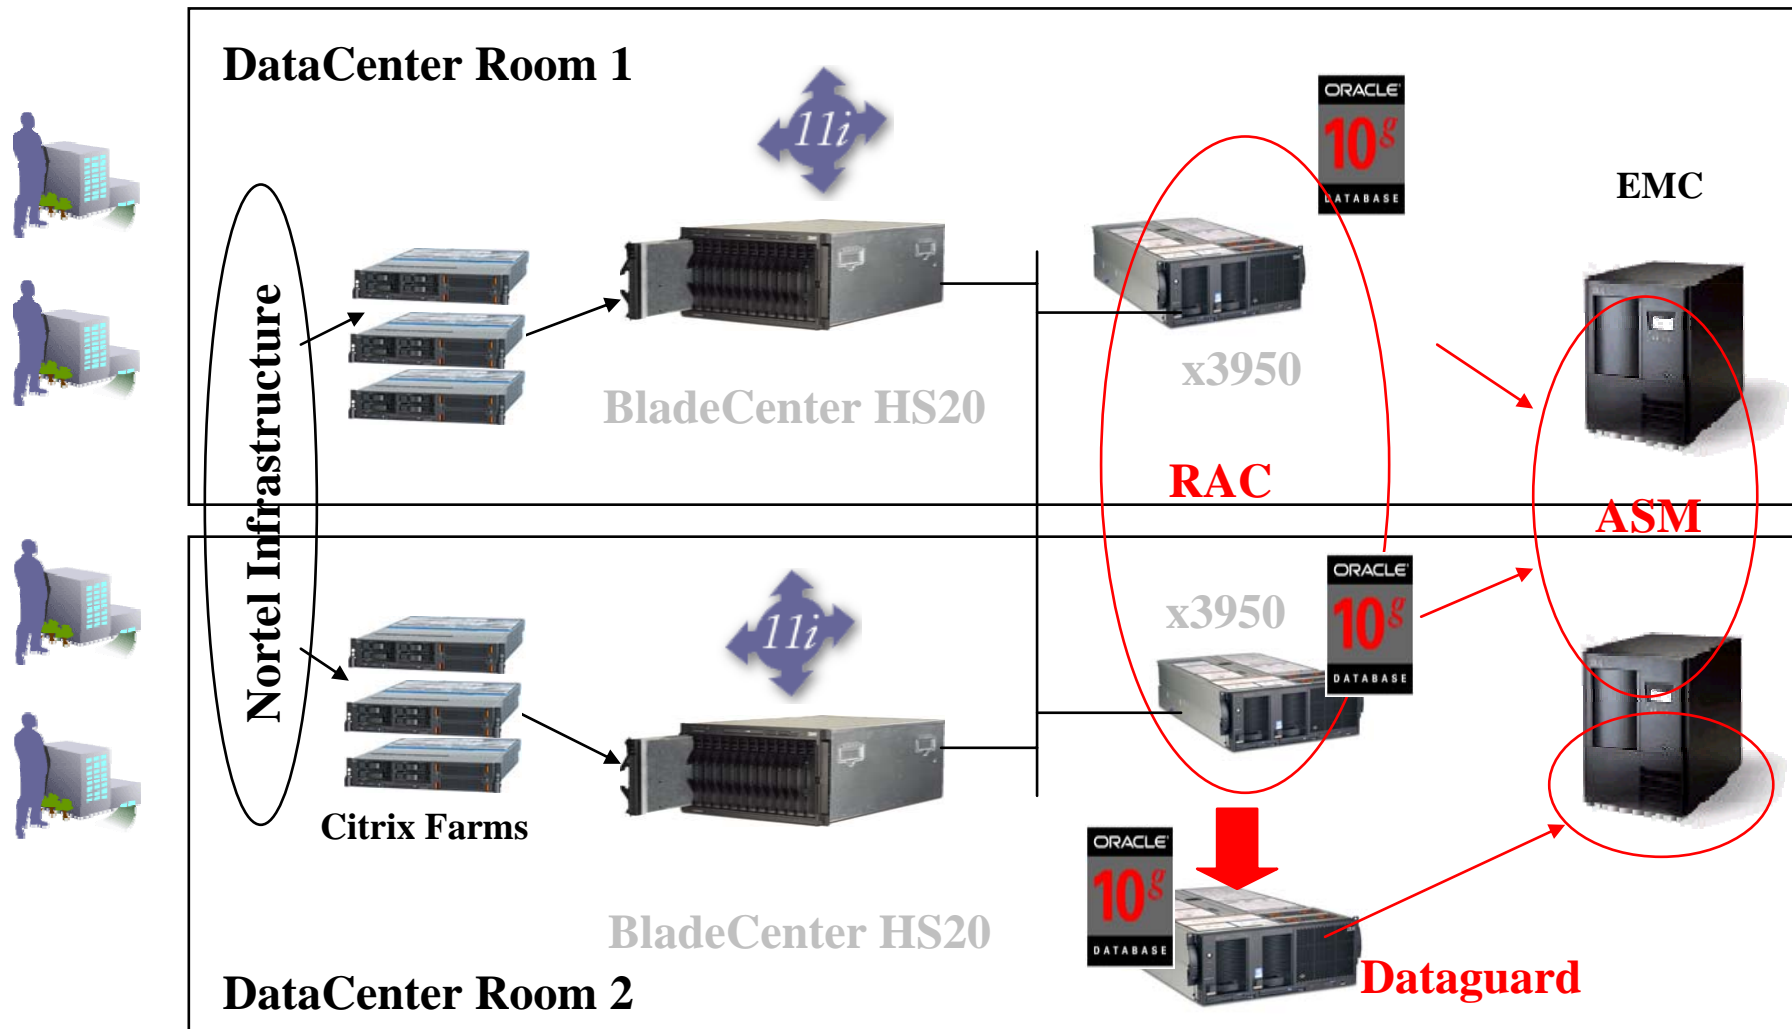
Yesterday IBM and Oracle announced a new initiative to grow Linux workload on the mainframe. Oracle has now certified PeopleSoft, Siebel, and eBusiness Suite to run with the backend datasever running on Linux on System z, and we anticipate those applications to be popular with enterprise customers. We're making a number of joint investments to support customers who seek to deploy Oracle solutions, including competency centers worldwide, *a joint solution center in Montpellier, France*, and the IBM Oracle International Competency Center in Silicon Valley.

Oracle OpenWorld 2006

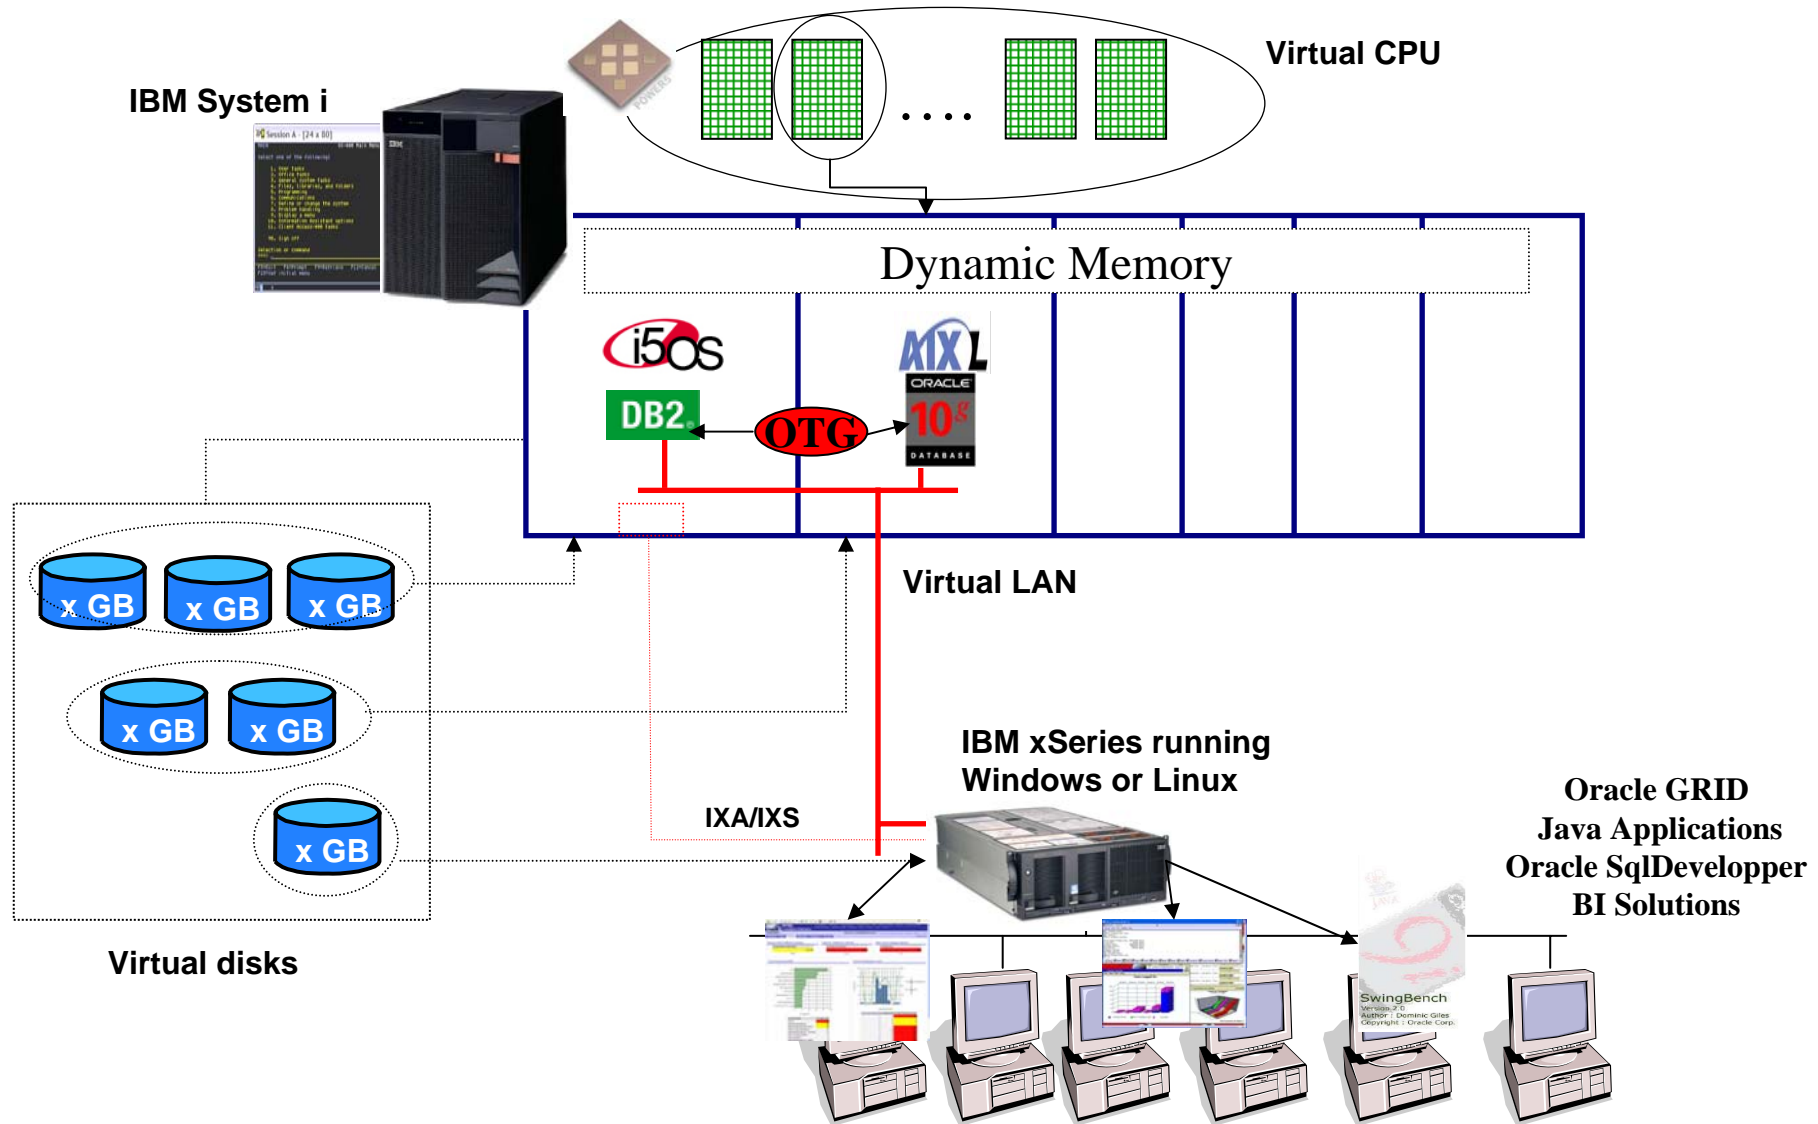
Oracle implementation on a IBM System z running Linux



11i/RAC project on IBM System x running Linux



Integrated New Oracle workloads Within a IBM System i



Merging GRID and VIRTUALIZATION ?

Scalability:


- System and performance expandability; ability to handle increased
- user and workload requirements without significant disruption

Scale Up: -> View as IBM SMP strategy

- Increasing the compute power of an SMP server by adding CPUs and memory to the same server;
- also called vertical scaling

Scale Out: -> View as ORACLE GRID strategy & RAC?

- Obtaining additional compute power by adding more servers so that the workload is spread across
- multiple systems; also called horizontal scaling
- **Grid is not RAC!** Grid is about optimisation of Ressources Utilization & Manageability so Oracle's Grid & IBM Virtualization are not opposite.
- **Oracle RAC** is about :
 - Availability
 - Scalability
 - Workload Isolation



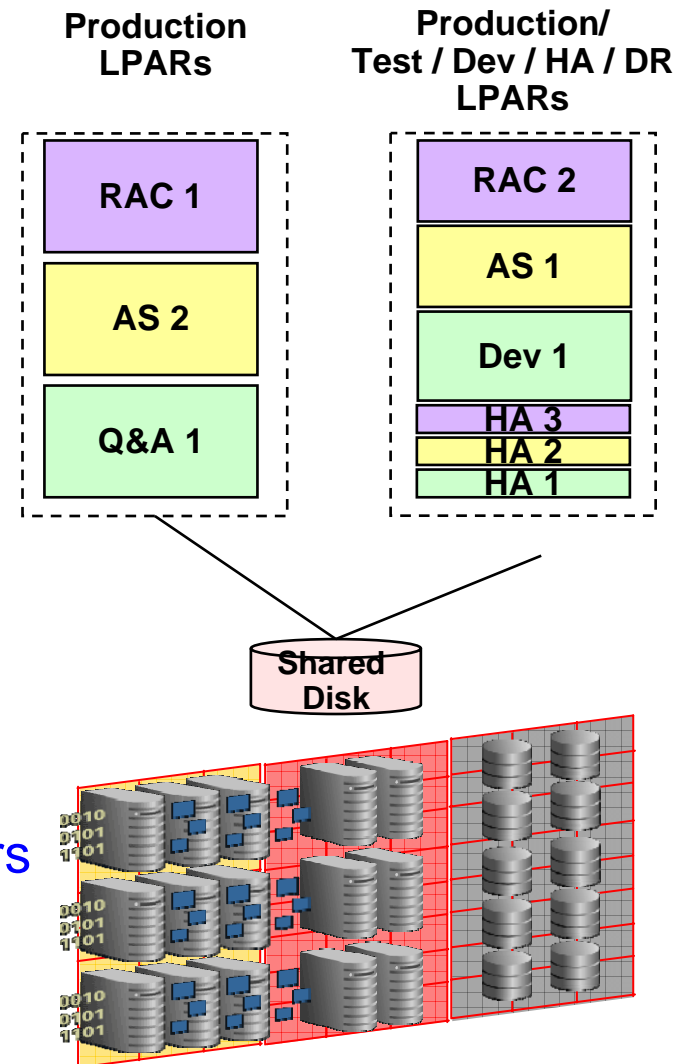
Merging GRID and VIRTUALIZATION ?

1. Scale up and Scale Out Solutions:
 - Both can be viable architectural choices
 - Depends upon different deployment requirements

2. Customer decisions are influenced by:
 - Application design, OS choice
 - Availability, Scalability, DR Approach
 - System Management
 - Workload Stability (Peaks, Batch), Affinity
 - Service Level Agreements
 - H/W acquisition cost vs. TCO (Manageability)

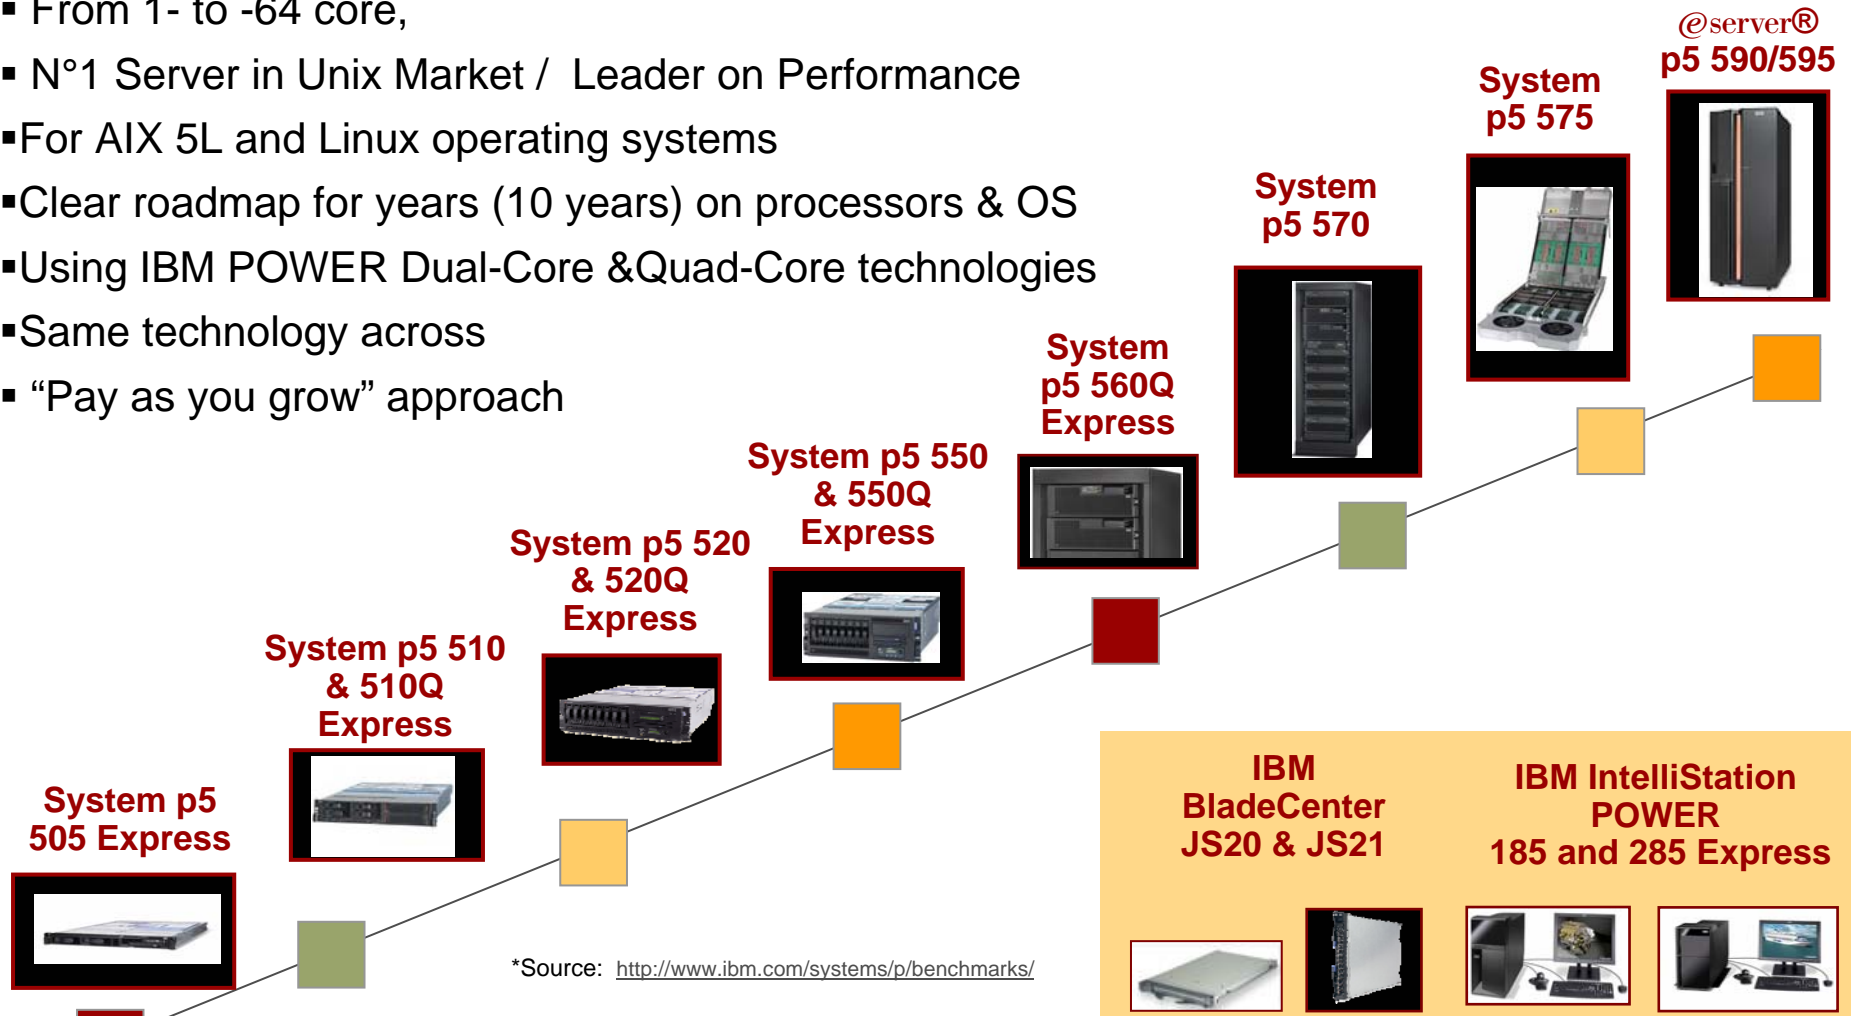
Gartner: Customer application deployments vary, The 'right answer' for the majority of customers is a combination of scale up and scale out... even within the same application deployment

Scale Up, Scale Out? Scale Right!



IBM Power5 – Scale Up, Scale Out, Scale within

- From 1- to -64 core,
- N°1 Server in Unix Market / Leader on Performance
- For AIX 5L and Linux operating systems
- Clear roadmap for years (10 years) on processors & OS
- Using IBM POWER Dual-Core & Quad-Core technologies
- Same technology across
- “Pay as you grow” approach




Most important Oracle/IBM Technology stacks used...

IBM *e*server

- IBM POWER5
- IBM Storage
- Virtualization
- SMT
- AIX5L/Linux
- GPFS
- HACMP

ORACLE[®] DATABASE 10^g



- Oracle Database
- 10g Application Server
- Grid Control
- RAC
- DataGuard
- ASM
- Transparent Gateways



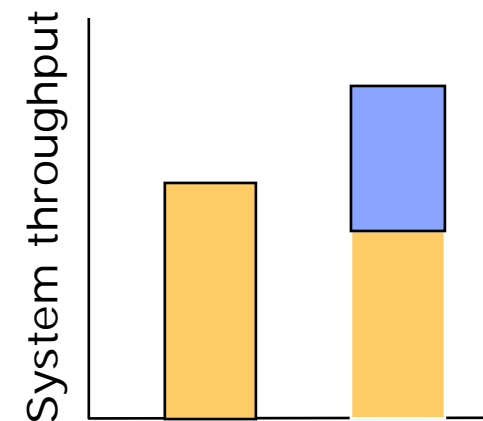
...to build joint projects

Simultaneous MultiThreading (SMT)

POWER5 (simultaneous multithreading)



Appears as 4 CPUs per chip to the operating system (AIX 5L V5.3 and Linux)

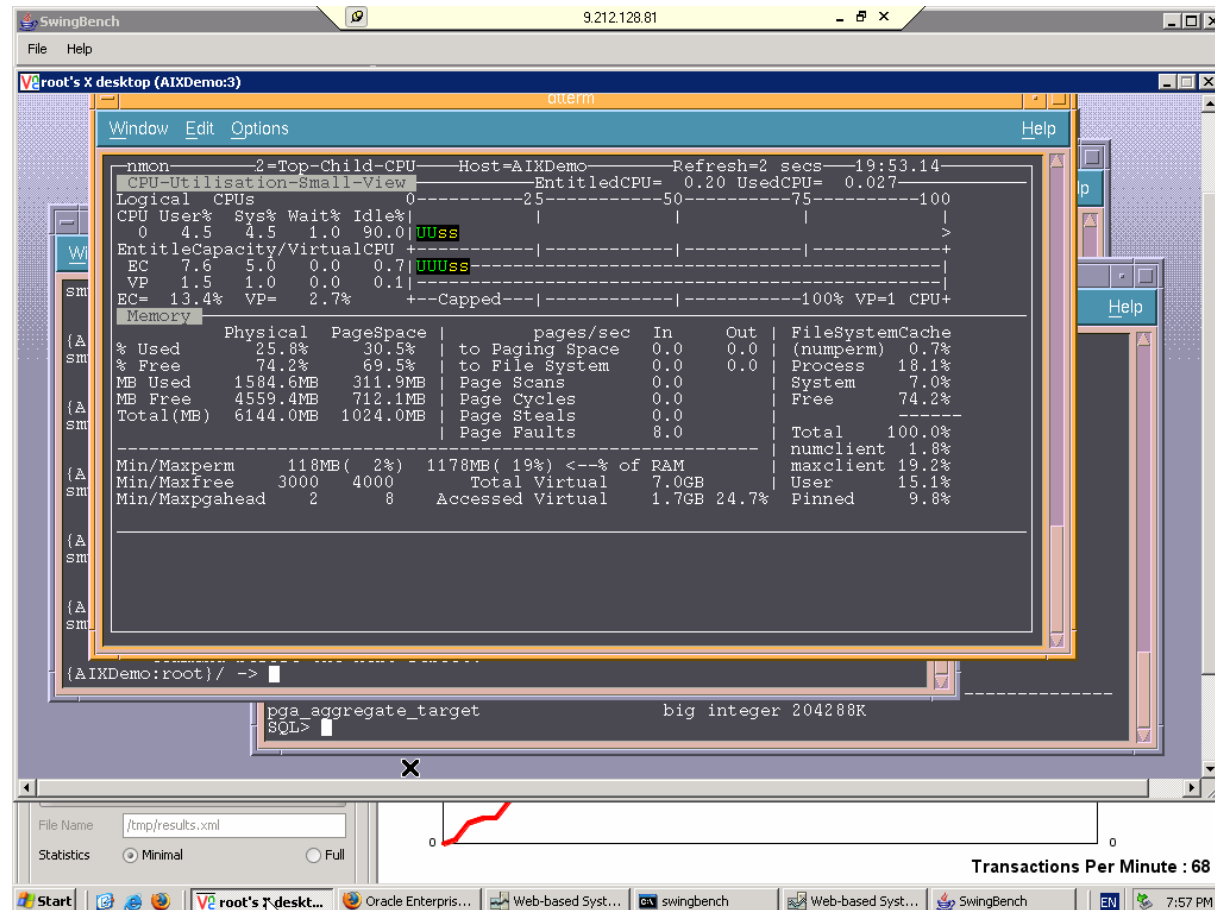


- Utilizes unused execution unit cycles – memory wait idle cycles
- Presents symmetric multiprocessing (SMP) programming model to software
- Natural fit with superscalar out-of-order execution core
- Dispatch two threads per processor: *“It’s like doubling the number of processors.”*
- Net result:

Power5 advanced virtualization options

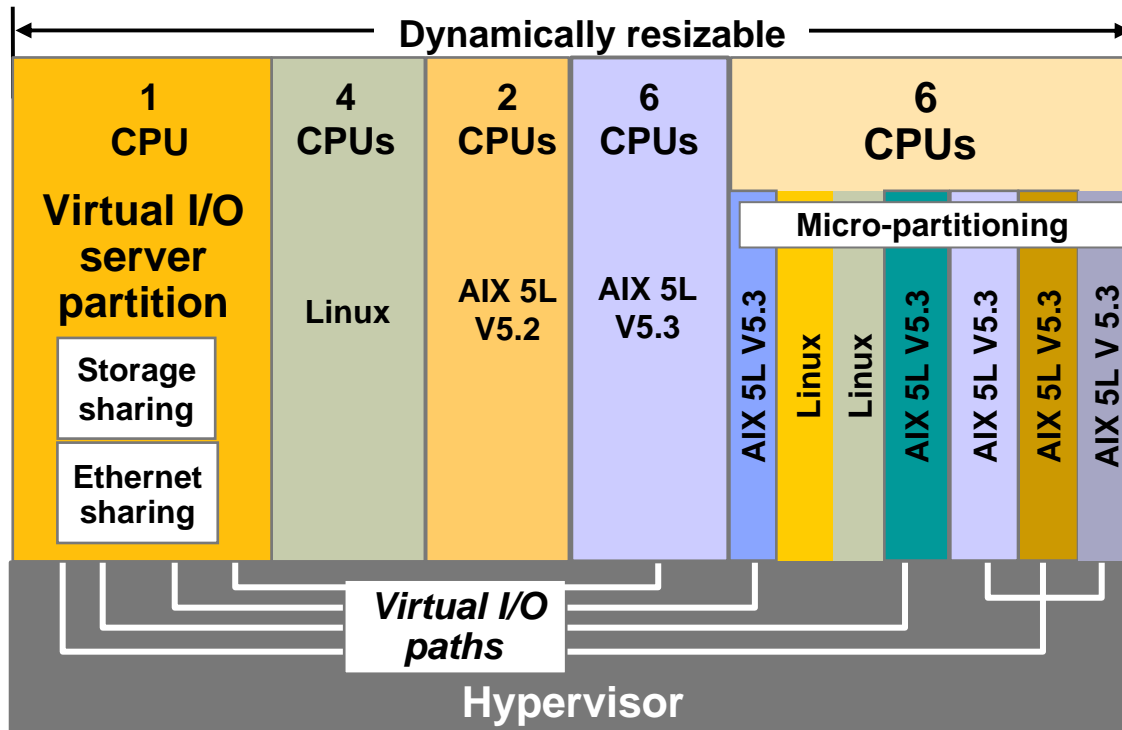
Optimize CPU in an Oracle/Power 5 environment

- **Activate SMT**
- Increase the Capacity Entitlement
- Add a Virtual Processor



- Optimize your Oracle/IBM environment to maximize your performance

IBM Power5 & AIX5L Advanced Virtualization



Virtual I/O server

- Shared Ethernet
- Shared SCSI and Fibre Channel-attached disk subsystems
- Supports AIX 5L V5.3 and Linux* partitions

Micro-Partitioning

- Share processors across multiple partitions
- Minimum partition 1/10th processor
- AIX 5L V5.3 or Linux*

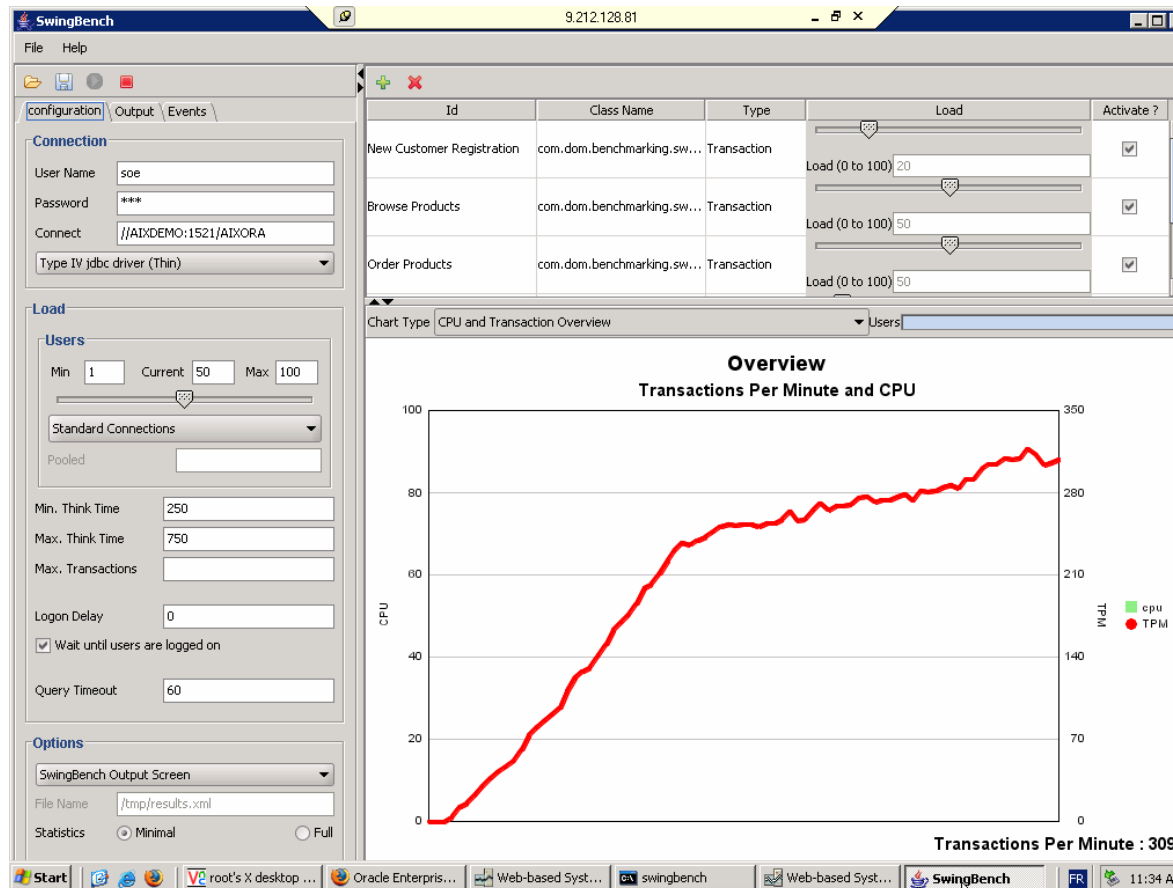
Partition Load Manager

- Balances processor and memory request

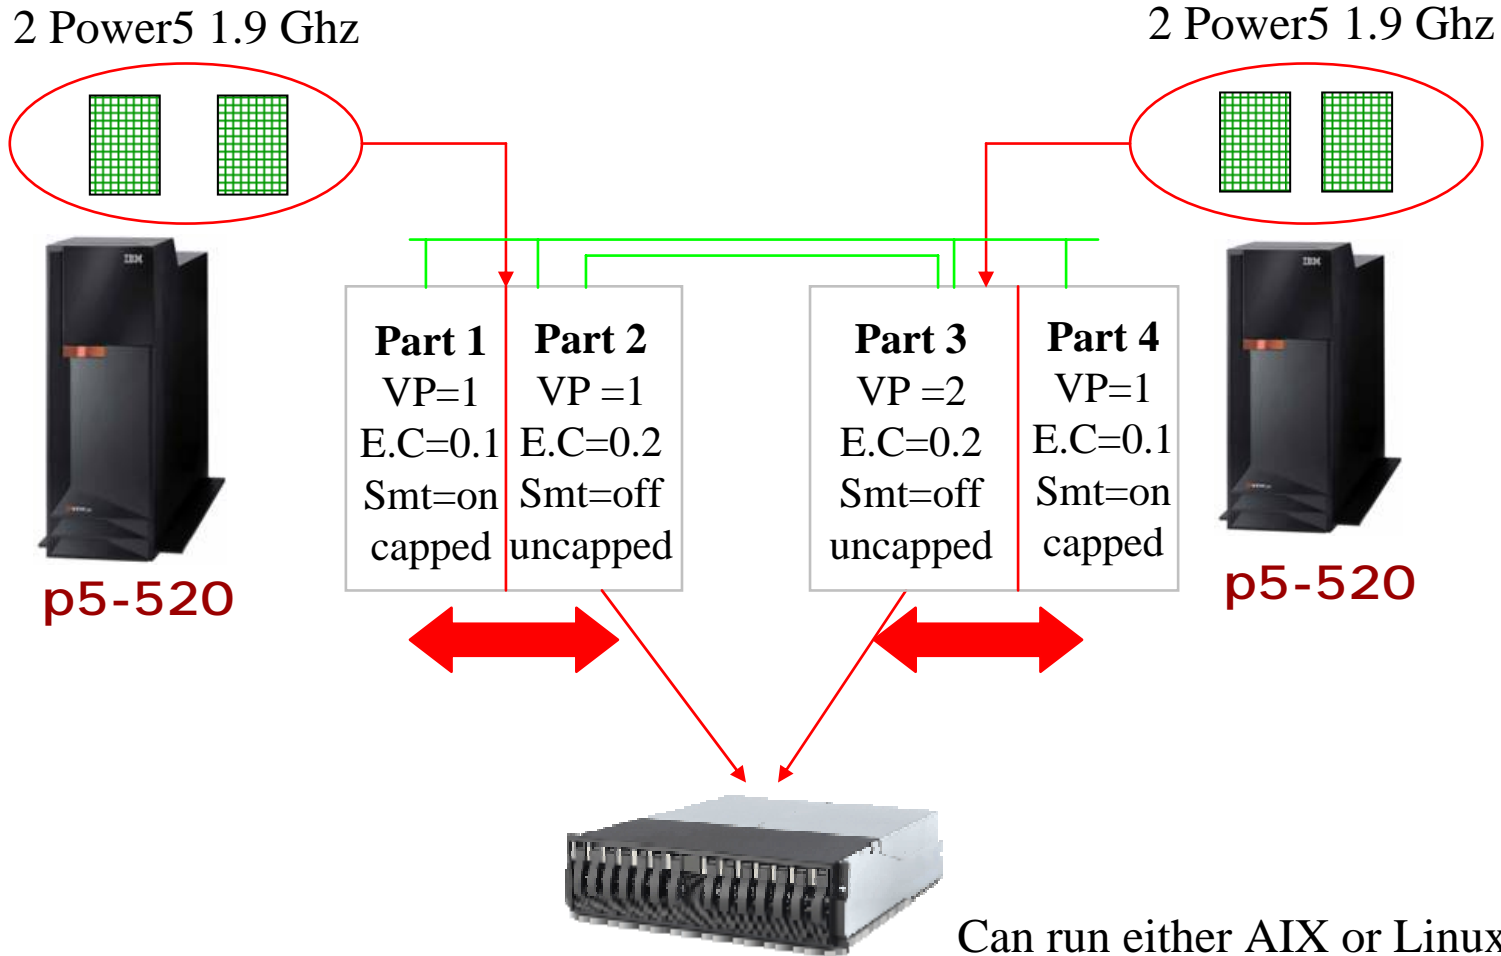


- DLPAR improvement with Oracle10g → automatic adaptation to CPU changes
- VIO Server supported for virtual network interface and disks for binaries
- IBM and Oracle are currently working together to confirm support for VIOS (disks) in a RAC environment. Certify will be updated when this is confirmed.
- 3rd Voting on NFS already certified on Linux x86 will be soon on AIX5L (JSC effort)

Power5 advanced virtualization options (Add Memory to Oracle on an IBM Power 5)

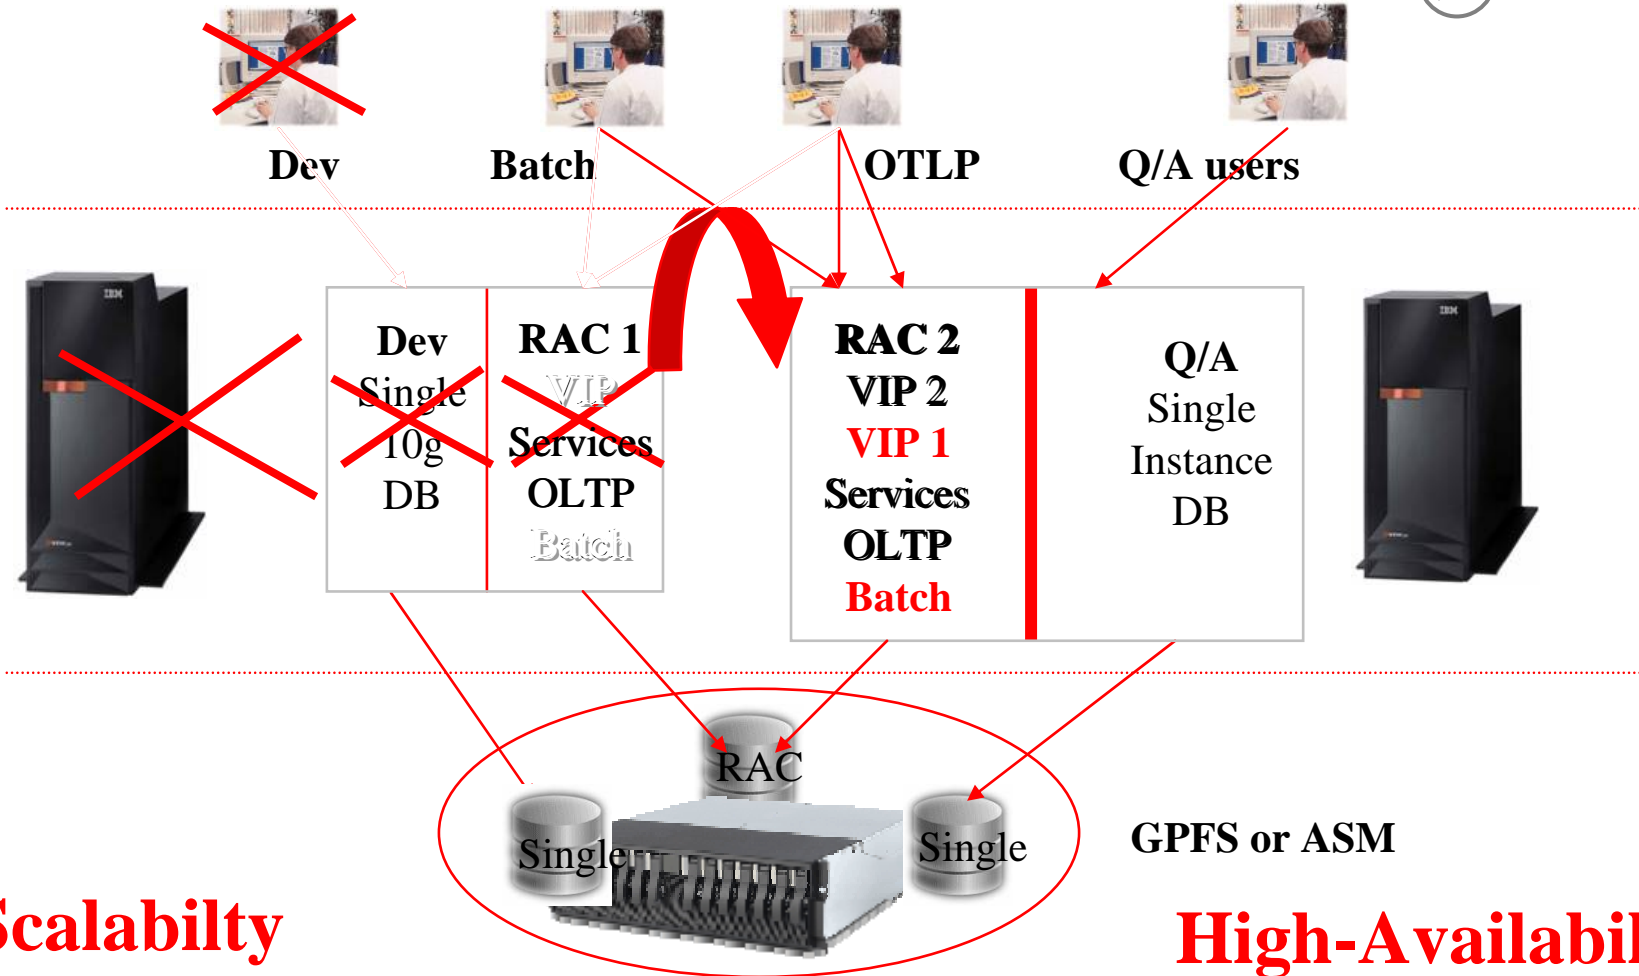


Get the best of your Scale I(*BM*)n & O(*racle*)ut solution



Get the best of your Scale I(BM)n & O(racle)ut solution

Virtualization



IBM General Parallel File-system (GPFS)

IBM filesystem designed for cluster use, certified with Oracle RAC

GPFS is capable of holding the database, redo logs, control files and archived logs required by Oracle RAC + Binaries

Data is striped across shared local disks or NSD servers

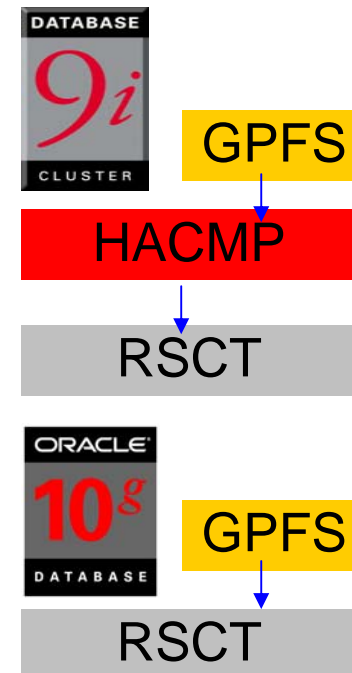
Metadata is maintained by all servers in the cluster
File Locking is distributed across the servers in the cluster

Performing and scalable

Striping I/O performance
Direct I/O with close to RAW device performance
Supports terabyte file sizes and multiple terabyte file systems

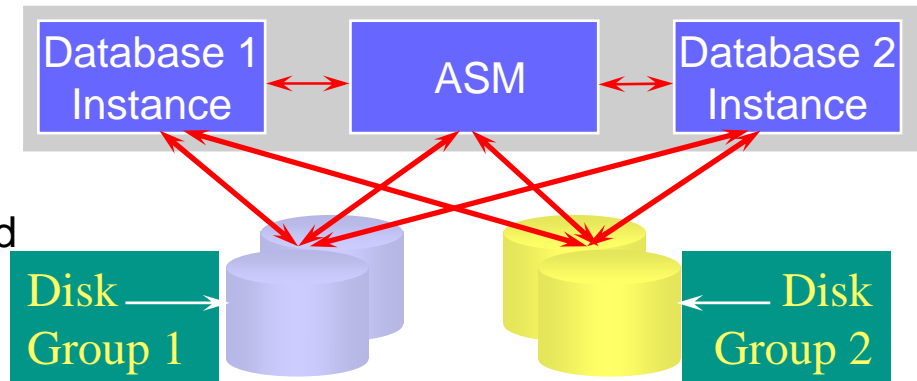
Highly Available – Fault tolerant

Supports RAID subsystems and/or performs own file system replication.
Supports on-line dynamic reconfiguration (add, delete, replace disks and nodes; rebalance the file system)

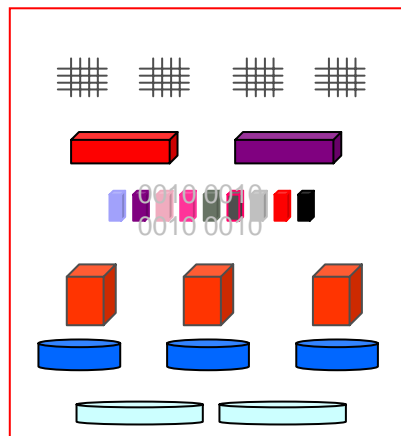


Storage Grid with Automatic Storage Management (ASM)

- Eliminates need for conventional file system and volume manager
- Capacity on demand
 - Add/drop disks online
- Automatic I/O load balancing
 - Stripes data across disks to balance load
- Automatic mirroring
- ASM extends SAME (Stripe and Mirror Everything)

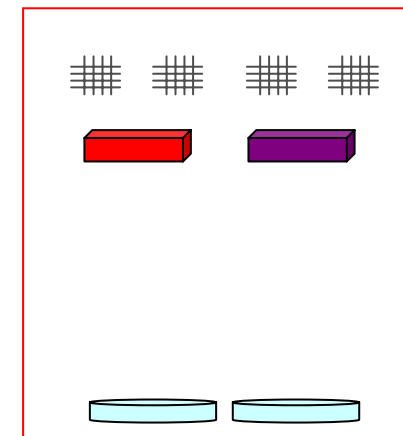


Non ASM



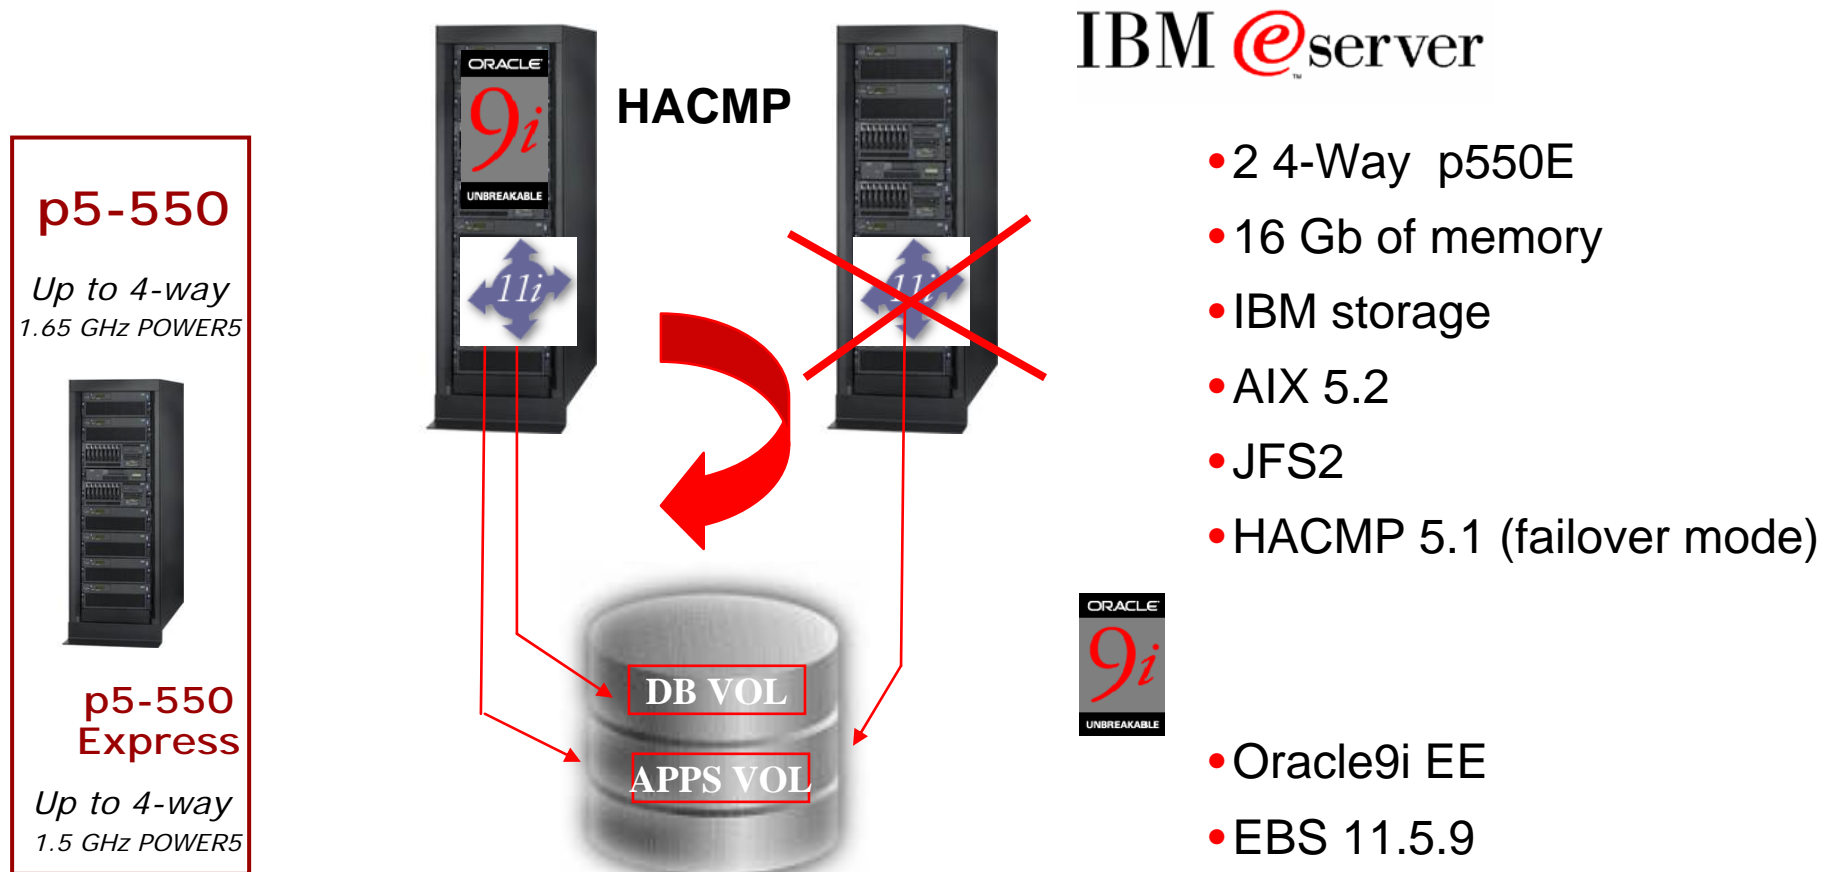
- ← Tables →
- ← Tablespace →
- ← Files →
- ← FileSystems →
- ← LVs →
- ← Disk Groups →

ASM



Oracle EBS 11i on a HACMP/p550 architecture

Original Low Cost Secured cluster solution



Oracle10g Standard Edition on System p


Low Cost Oracle RAC cluster solution

Customer plans to implement 11 distributed sites.

p5-520
Up to 2-way
1.65 GHz POWER5



p5-520 Express
Up to 2-way
1.5 GHz POWER5



IBM @server

All the benefits of the Oracle & IBM technologies on a single core 5K\$ servers !



Can run either AIX or Linux on Power

« This presentation is for informational purposes only and may not be incorporated into a contract or agreement. »


Oracle9i RAC and EBS 11i on a Dynamic p570

RAC in a « Micropartitoned » BOX

p5-570
 Up to 16-way
 1.65 GHz POWER5
 1.9GHz POWER5



Up to 8-way
 1.5 GHz POWER5



IBM server

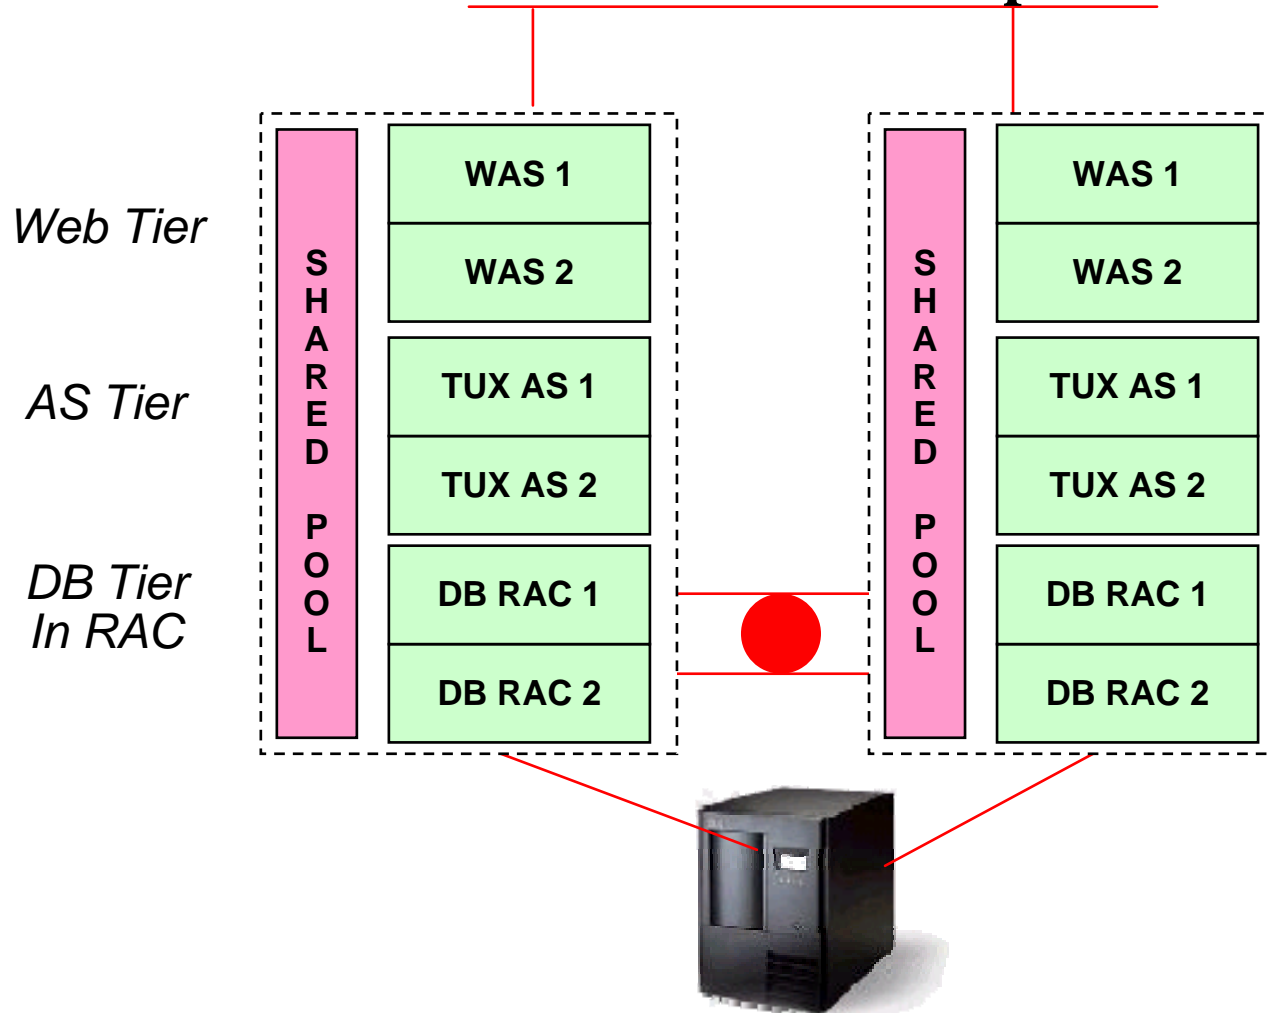
- 1 16-Way p570
- 2 DLPAR
- 16 Gb of memory
- Micro-partitioning
- 1 ESS800
- AIX 5.3 ml02

ORACLE DATABASE

- Oracle9i EE
- EBS 11.5.8
- RAC
- GPFS

Scale Strategy with Oracle GRID & IBM virtualization

A scale « in-out » ERP implementation



IBM *e*server

- 3 16-Way p570
- 2 DS6800
- AIX 5.3 ml04
- GPFS

ORACLE DATABASE 10^g

- Oracle10g EE
- RAC
- PeopleSoft CRM

Optimized CPU and memory resources between all logical tiers

Production Real Customer Case -> Afternoon

