APV In Practise

Live Production Environments and Benchmarks Experiences with Advanced POWER Virtualisation

Jean-Armand Broyelle
jabroyelle@fr.ibm.com
PSSC System p benchmark

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, BladeCenter, Blue Gene, Blue Lightning, C Set++, CICS, CICS/6000, ClusterProven, 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, Passport Advantage, POWERparallel, PowerPC, PowerPC(logo), Predictive Failure Analysis, PS/2, pSeries, PTX, ptx/ADMIN, RETAIN, RISC System/6000, RS/6000, RT Personal Computer, 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, TrackPoint, TURBOWAYS, UltraNav, 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, Chipkill, Cloudscape, DB2 OLAP Server, DB2 Universal Database, DFDSM, DFSORT, Domino, e-business(logo), e-business on demand, eServer, Express Middleware, Express Portfolio, Express Servers, Express Servers and Storage, GigaProcessor, HACMP, HACMP/6000, Hypervisor, i5/OS, IBMLink, IMS, Intelligent Miner, Micro-Partitioning, iSeries, NUMACenter, ON DEMAND BUSINESS logo, 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, System p5, S/390 Parallel Enterprise Server, ThinkVision, Tivoli Enterprise, TME 10, TotalStorage Proven, Ultramedia, VideoCharger, Virtualization Engine, 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 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 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 July 20, 2005

Notes on benchmarks and values

The IBM benchmarks results shown herein were derived using particular, well configured, development-level and generally-available computer systems. 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 Web site of the benchmark consortium or benchmark vendor.

IBM benchmark results can be found in the IBM System p5, @server p5, pSeries, OpenPower and IBM RS/6000 Performance Report at http://www.ibm.com/servers/eserver/pseries/hardware/system_perf.html.

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.

For a definition and explanation of each benchmark and the full list of detailed results, visit the Web site of the benchmark consortium or benchmark vendor.

TPC	http://www.tpc.org
SPEC	http://www.spec.org
LINPACK	http://www.netlib.org/benchmark/performance.pdf
Pro/E	http://www.proe.com
GPC	http://www.spec.org/gpc
NotesBench	http://www.notesbench.org
VolanoMark	http://www.volano.com
STREAM	http://www.cs.virginia.edu/stream/
SAP	http://www.sap.com/benchmark/
Oracle Applications	http://www.oracle.com/apps_benchmark/
PeopleSoft - To get information on PeopleSoft benchmarks, contact PeopleSoft directly	
Siebel	http://www.siebel.com/crm/performance_benchmark/index.shtm
Baan	http://www.ssaglobal.com
Microsoft Exchange	http://www.microsoft.com/exchange/evaluation/performance/default.asp
Veritest	http://www.veritest.com/clients/reports
Fluent	http://www.fluent.com/software/fluent/fl5bench/fullres.htmn
TOP500 Supercomputers	http://www.top500.org/
Ideas International	http://www.idesinternational.com/benchmark/bench.html
Storage Performance Council	http://www.storageperformance.org/results

Revised July 5, 2005

Notes on performance estimates

rPerf

rPerf (Relative Performance) is an estimate of commercial processing performance relative to other IBM UNIX 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 August 12, 2005

Introduction

- What is this presentation about ?
  - Many IBM customers have already switched to Advanced POWER Virtualisation in their production environment. Several benchmarks and tests have also been performed involving mixed applications solutions using SPLPAR and VIOS.
  - We will share results and experiences from benchmarks and describe the new architecture capabilities addressed by this technology through some real-life implementations. We will also emphasize trends/ lessons from these cases.
- Audience:
  - Architects and system administrators
  - Good knowledge of POWER5 virtualization features required

Agenda

- APV technology reminder
- Sizing and Benchmark results with APV solutions
- Customer and partner Production experiences
  - Consolidation examples
  - Ephemeral environments and provisioning support
  - HA architecture, APV and CoD combinations

Agenda

- APV technology reminder
- Sizing and Benchmark results with APV solutions
- Customer and partner Production experiences
  - Consolidation example
  - Ephemeral environments/provisioning support
  - HA architecture, APV and CoD combinations

@server p5 Advanced Virtualization Option : Reminder

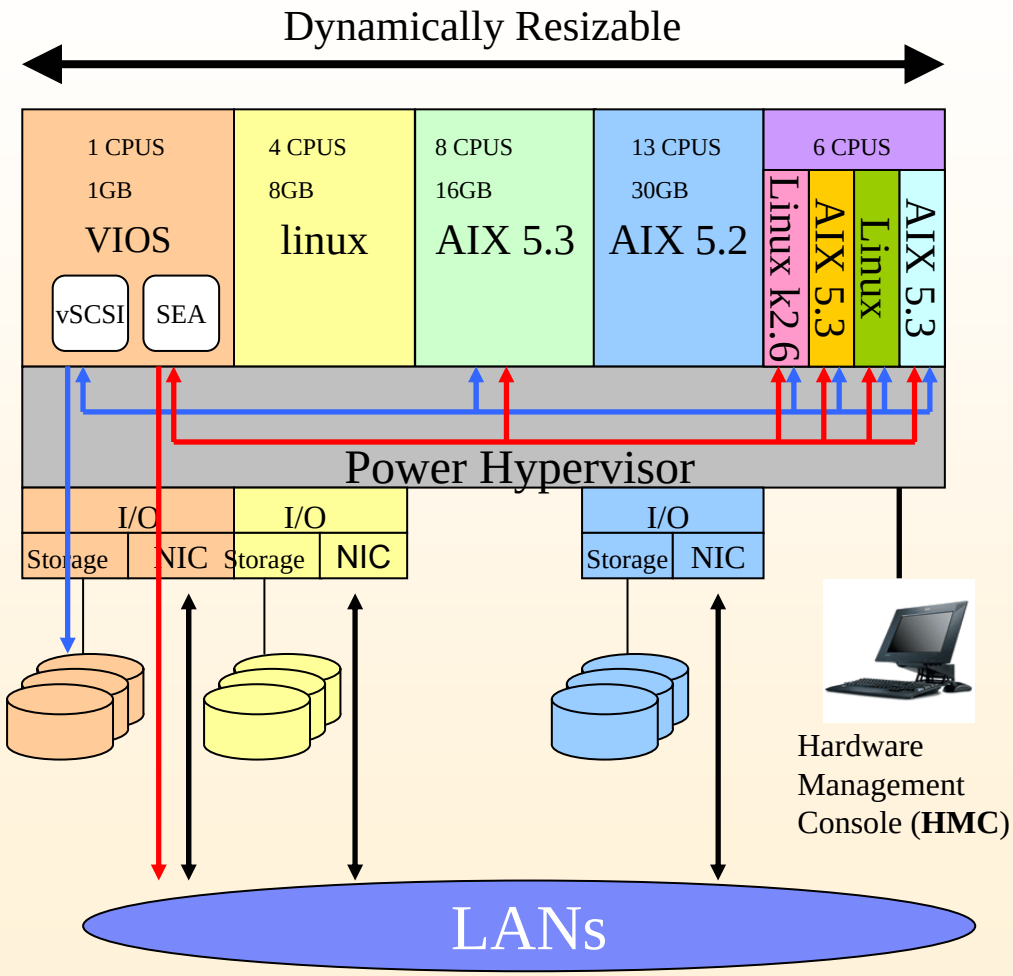

Increase Physical resource utilization thru virtualization of processors, memory, network and disk ressources

Micro-partition (Shared Processor LPAR)

- A single processor might be shared by up to 10 partitions : Support for up to 254 partitions
- SPLPAR attributes: CE, VPs, Capped/uncapped mode, uncapped weight

Virtual inter-partition Ethernet

- Virtual Ethernet-LPARs can communicate without having to use a physical I/O adapter
- Ethernet Sharing (SEA) – LPARs can share external network connection

Virtual I/O (disk)

- Client partitions can use logical disks hosted by another partition

Agenda

- APV technology reminder
- Sizing and Benchmark results with APV solutions
- Customer and partner Production experiences
  - Consolidation example
  - Ephemeral environments/provisioning support
  - HA architecture, APV and CoD combinations

Sizing and Benchmark results with APV solutions

- Why a benchmark ?
  - To measure the optimal performance of a solution regarding workload and hardware requirements
    - Way to know or prove a hardware sizing (major feed-back from benchmarks)
- Benchmarking: two possible targets
  - Application limit benchmark : the benchmark output is the maximal application performance we can get with the available hardware resources .
  - Platform sizing benchmark : the application performance target is known, the benchmark output is the suitable hardware configuration (at the lowest cost).

APV : performance dimensions

- POWER Virtualisation goal: optimize the physical ressource usage
  - Theoretically it should be good for performance
  - APV implementation is also resource consuming
- Possible causes
  - Hypervisor, efficiency of the physical CPU in the Shared processor pool, VIO memory and CPU consumption, SEA and vSCSI performances
- Tunables settings:
  - Capacity Entitlement, VPs number, capped/uncapped, shared processor number
  - SEA adapters configurations, VIOS Disk IOs adapters, CE and VPs of VIOS LPARs

Benchmark results with APV solutions

- No official or ISVs benchmarks, nor public results published on standard benchmark yet
- customer and ISV benchmarks stories performed at PSSC:
  - 10 tests performed in Montpellier involving APV features successfully :
    - APV resources Sizing,
    - optimize the hardware performance using APV features
  - 7 tests examples of using APV features performed in the EMEA PSSC pSeries and system p5 benchmark center.
  - feed-backs to share:
    - Get sizing rules for micro-partitions and VIOS **(platform limit benchmarks)**.
    - **Application tunings:** using APV helps the application to run faster with the same hardware.
  - Different solutions tested:
    - WebSphere, MQ Series
    - Oracle database (9i, 10g)
    - SAP application server, PeopleSoft, AMDOCS

WebSphere benchmark

- configuration :
  - Application server : 16 ways p570 1.9GHz , 128GB RAM
  - AIX 5.3 ML1, VIOS 1.1
  - 6 LPARs
    - 2 WEB seal,
    - 2 WEBSphere AppServer,
    - 1LDAP,
    - 2VIOS
- Goal : APV sizing
- Load Runner Transactional Workload
  - Same OLTP workflow : Synchronous CPU activity across all the architecture components

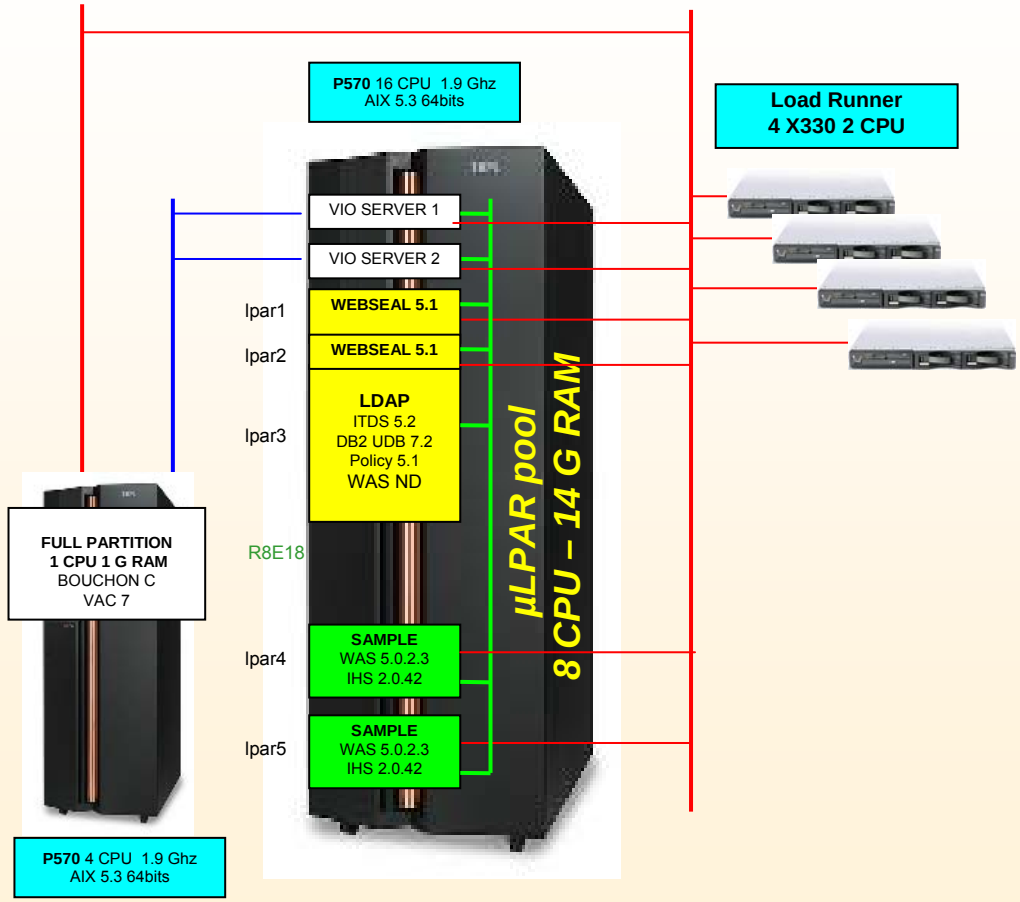

SPLPAR / Dedicated LPAR ratio : what to compare?

Measure the minimal CPU resources for same application performance criteria and quality of services in Dedicated and Shared mode.

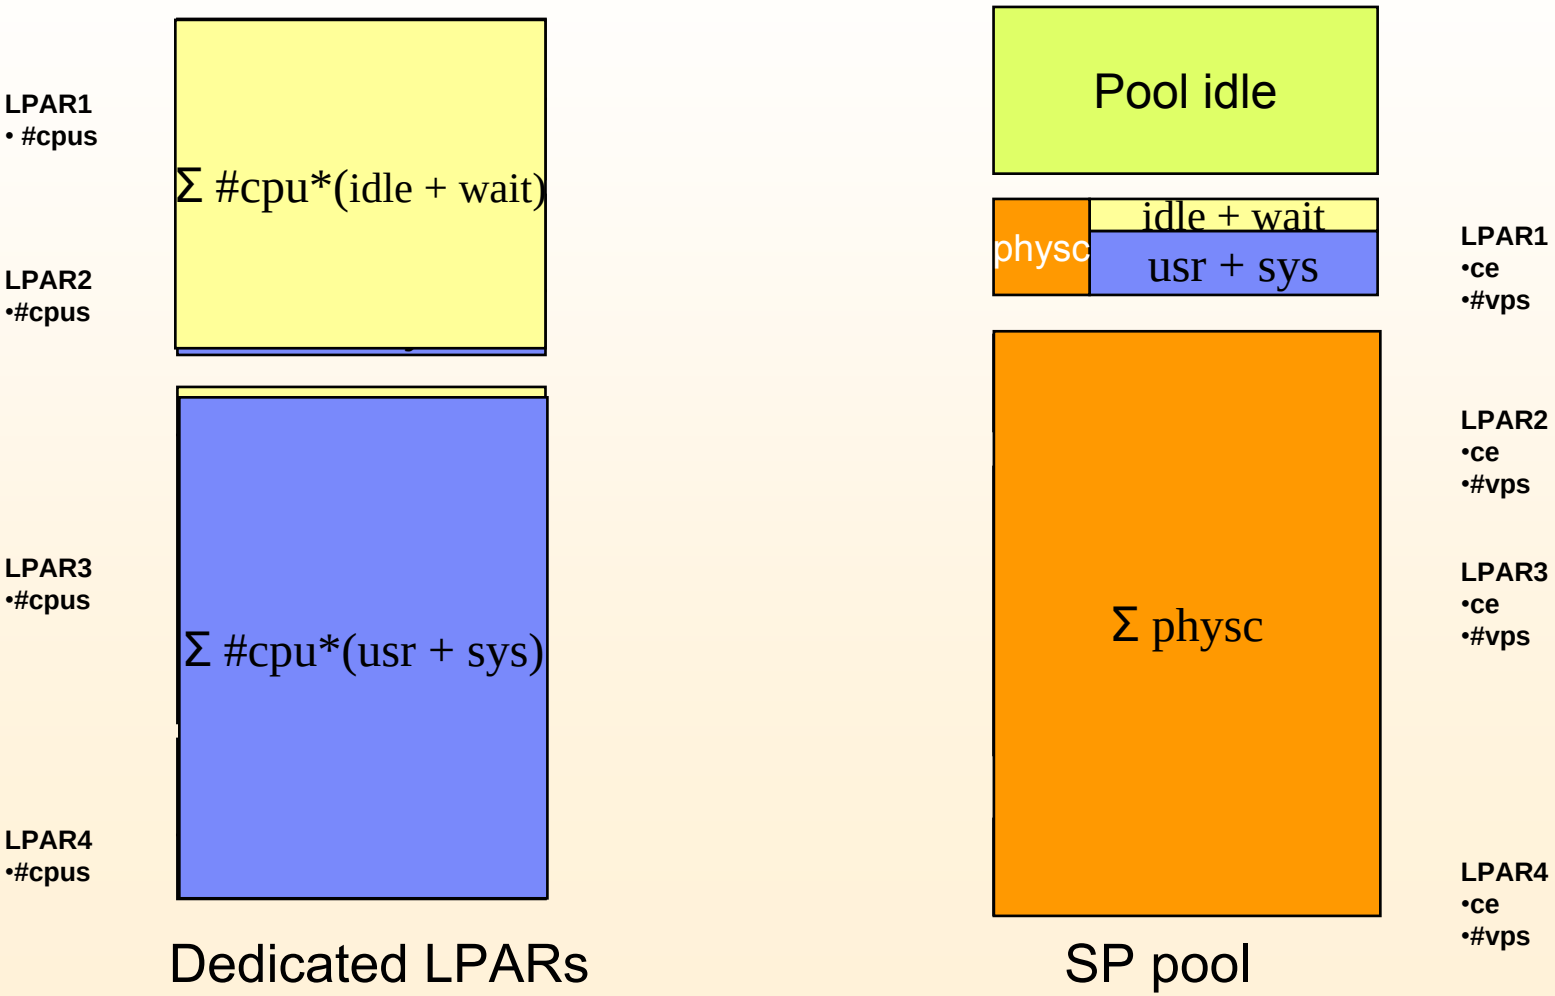

SPLPAR / Dedicated LPAR ratio:what to compare?

Answer : physical CPU number of the 2 configurations and then the ratio of CPU consumption (for sizing purpose)

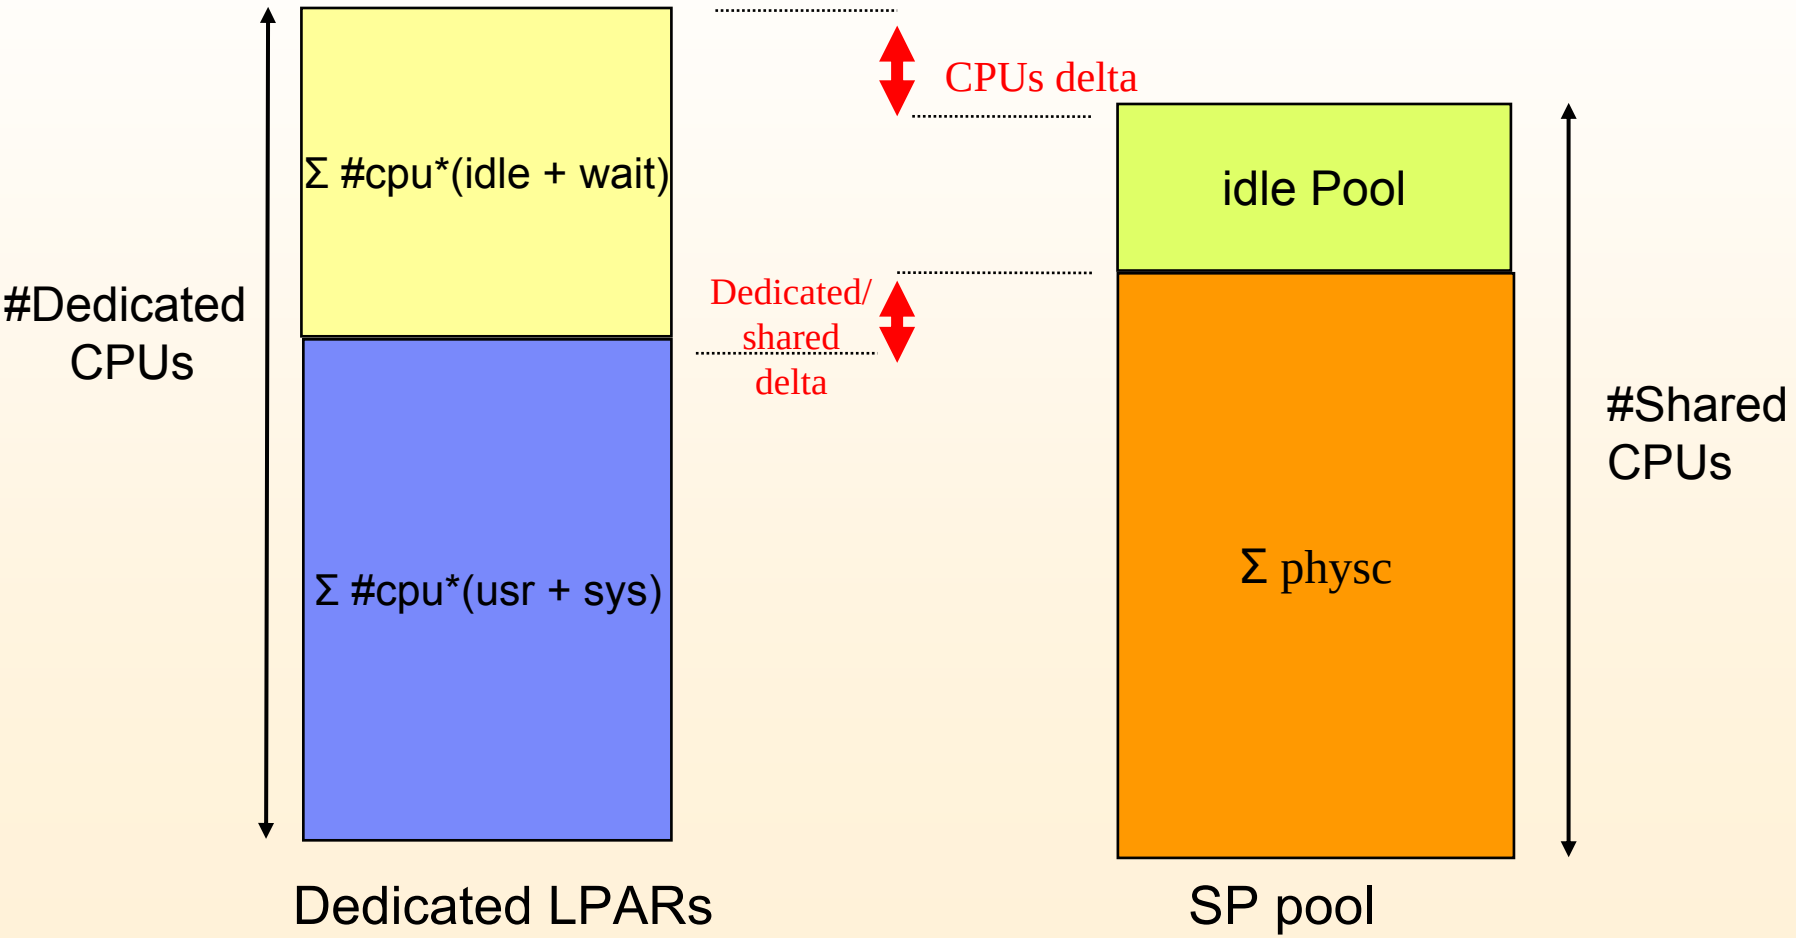

SPLPAR / Dedicated LPAR ratio

measure the Shared/Dedicated CPU ratio

- 2 measurements with the same application performance criteria and quality services
- Calculate the 2 following ratios
  - Processors configuration ratio (P)

P= #SharedCPUs $_{shared}$ / Σ $_{dlpars}$ #CPU $_{dedicated}$

  - CPU consumption ratio (C)

C= Σ $_{splpars}$ Physc$_{shared}$ / Σ $_{dlpars}$ (User%$_{dedicated}$ + Sys%$_{dedicated}$)*#CPU$_{dedicated}$

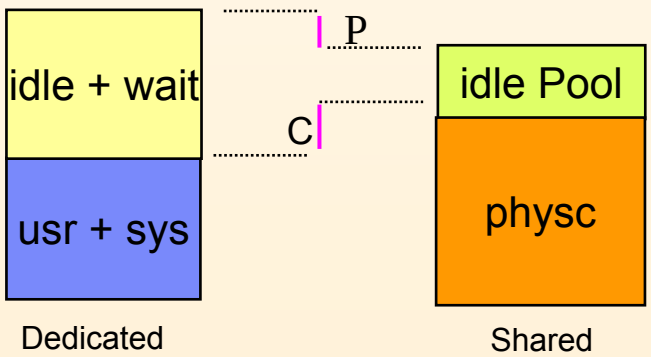

WebSphere benchmark

- Dedicated/Shared configurations comparison :
  - Fixed performance criteria : same LR configuration, same Transaction response time, same throughput

Dedicated :9 CPUs Dedicated

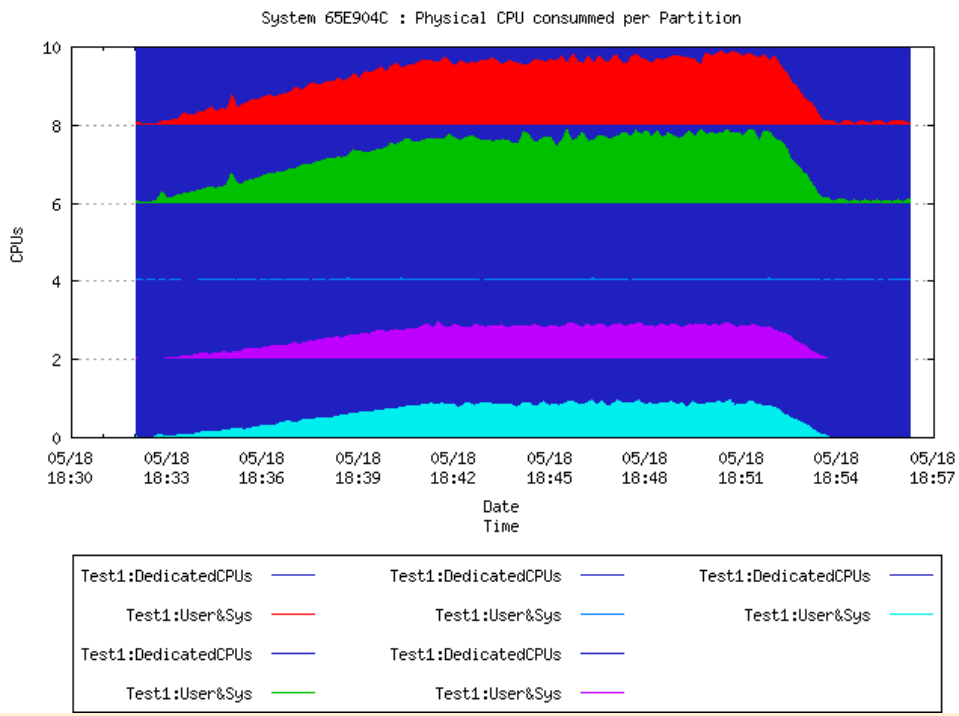

Shared : 7 shared processors, uncapped.

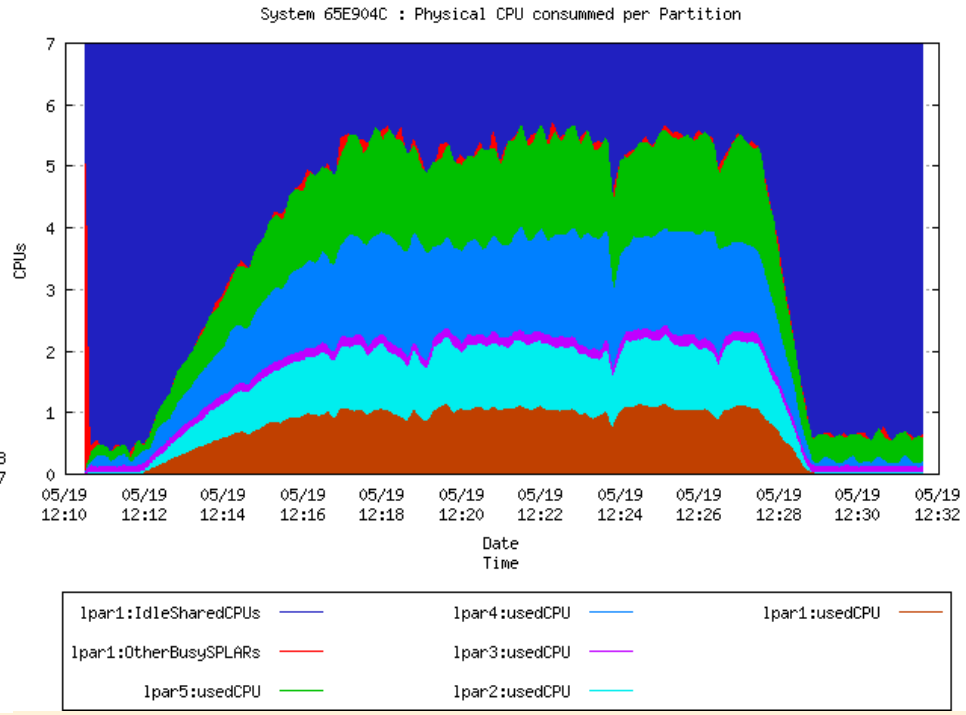

WebSphere benchmark

- Results :
  - Processors configuration
    - P =~ 77%
  - CPU consumption
    - C =~ 115%

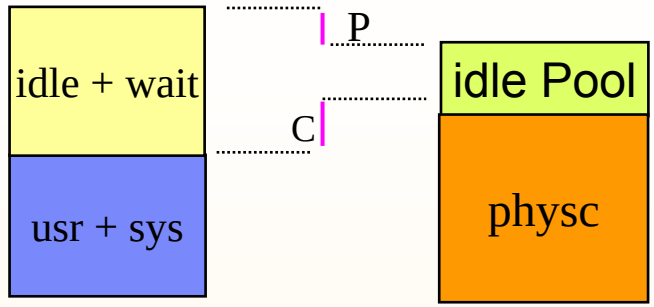

- Analysis :
  - CPU consumption on LPAR have same profile: Websphere and network load balancing. → CPUs number in the shared processor pool is close to Dedicated.
  - one partition is almost idle LDAP : Physc =~0.1 versus , 1 CPU in Dedicated → P=1000%.

PeopleSoft benchmark

- Configuration :
  - 3* p5 595 : 64@1.9GHz 256GB RAM
  - SunFire 25K
  - DS 8300
- Software stack
  - AIX 5.3 ML3, no VIO, GPFS
  - People Soft 8.46
  - Oracle RAC 10gR2
  - MQ 5.3 .05
  - WAS 5.1.1
  - Tuxedo
- Benchmark target: OLTP and Batch Stress Test

PeopleSoft benchmark

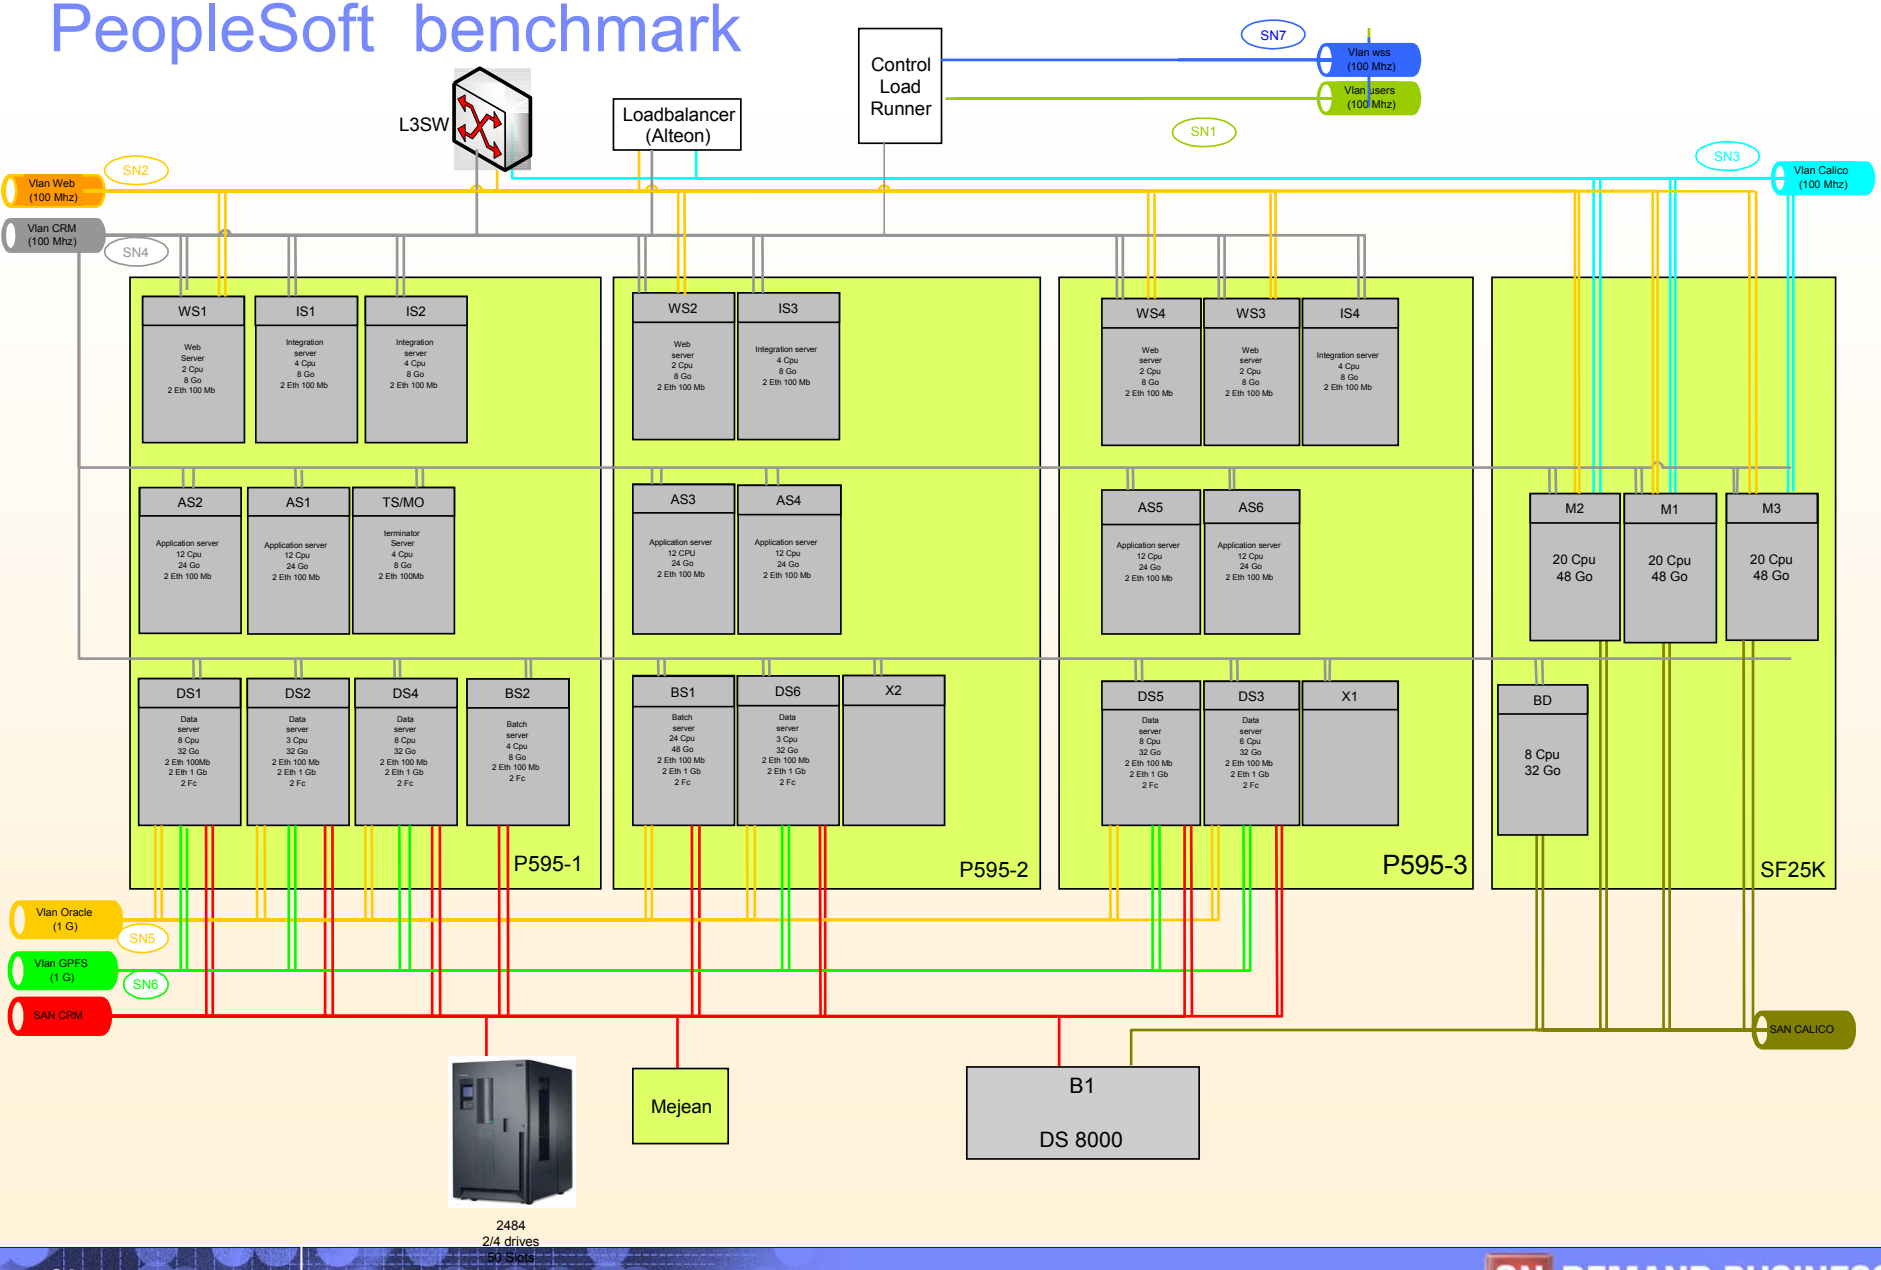

PeopleSoft benchmark:results

- Dedicated/Shared configurations comparison: 2,400 running virtual users

Dedicated mode :3 * 64 CPUs Dedicated
sustain Throughput :1,110,000 bytes/s
Action_Transaction: 59.07 sec

Shared mode : 3*64 shared processors,
sustain Throughput :1,112,000 bytes/s
Action_Transaction: 53.256 sec

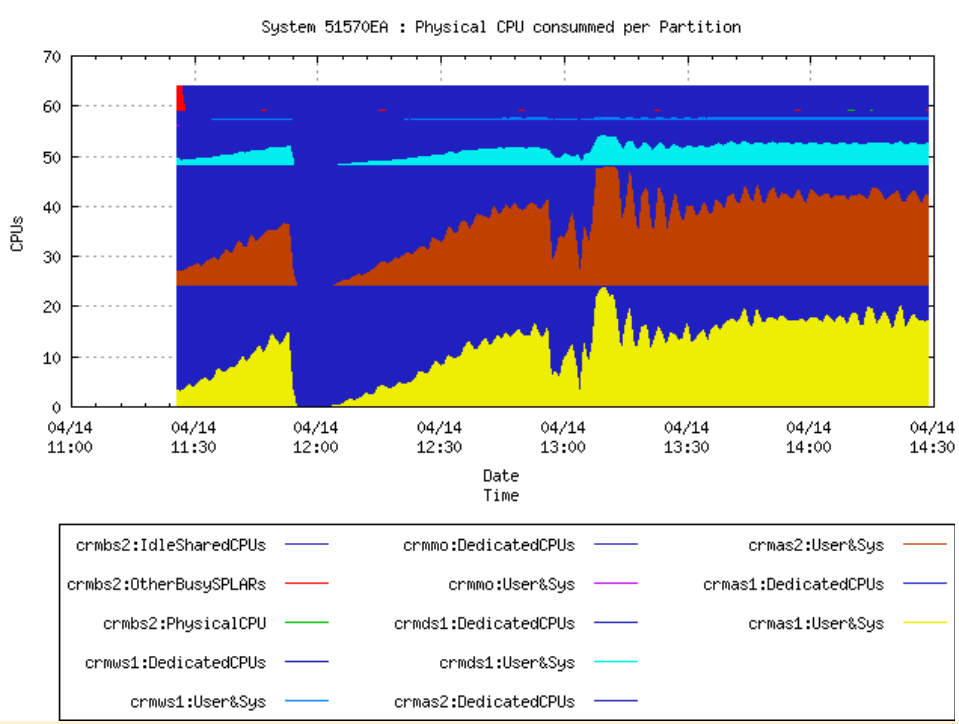

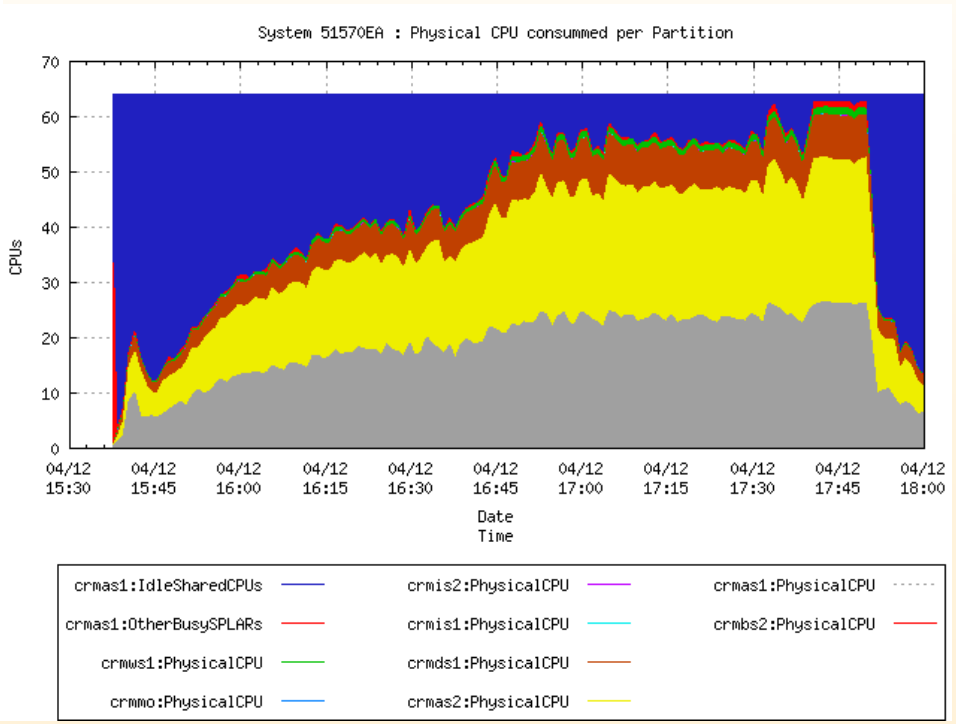

PeopleSoft benchmark: analysis

- Comparison Results :
  - Processors configuration (P) = 100%
  - CPU consumption (C) =~ 120%
  - Better Response time

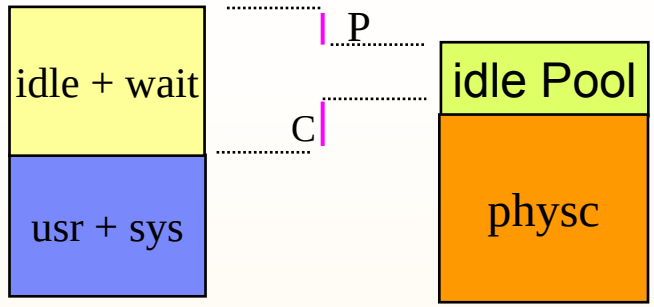

- Analysis :
  - static sizing for DLPARs mode (same CPUs number for all different test scenario) → LPARs CPU usage may vary from test to test. Ideal for a SPLPAR configuration.
  - Higher CPU consumption in shared mode because of VPs switches.
  - Performance highest reference reached with SPLPARs for all application scenario

Oracle 9i benchmark result

- Configuration:
  - Application server : 16 ways p570 1.9GHz , 128GB RAM
  - AIX 5.3 ML2, no VIOS
  - 3 LPARs : different CPU activity profiles
    - Oracle RAC DB : Oracle application 11i instance + BO instance
    - 1 Oracle Application
- Transactional Workload (Load Runner) : 2 workloads
  - Oracle Financial Tx
  - BO queries

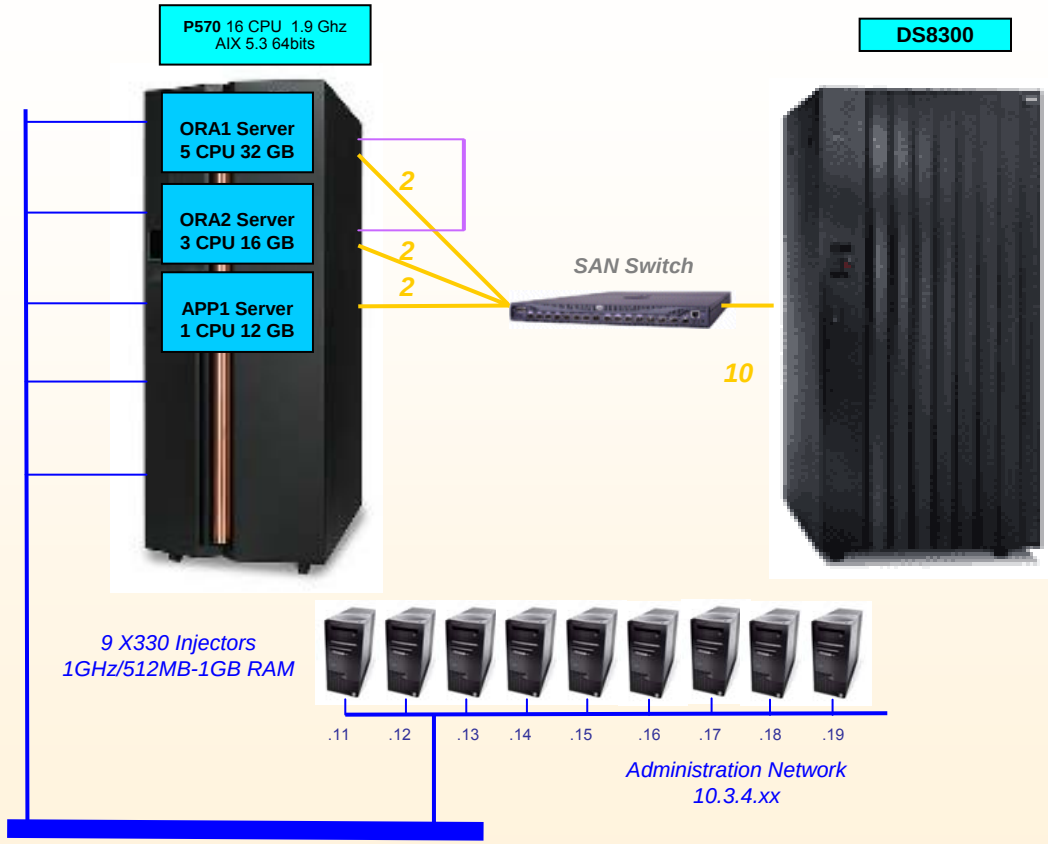

Oracle 9i benchmark: results

- Dedicated/Shared configurations comparison:
  - 333 Load Runner virtual users,
  - fixed throughput (← adjustable think time)

Dedicated :6 CPUs Dedicated

Transaction Response time (90%) : **13.25 sec**

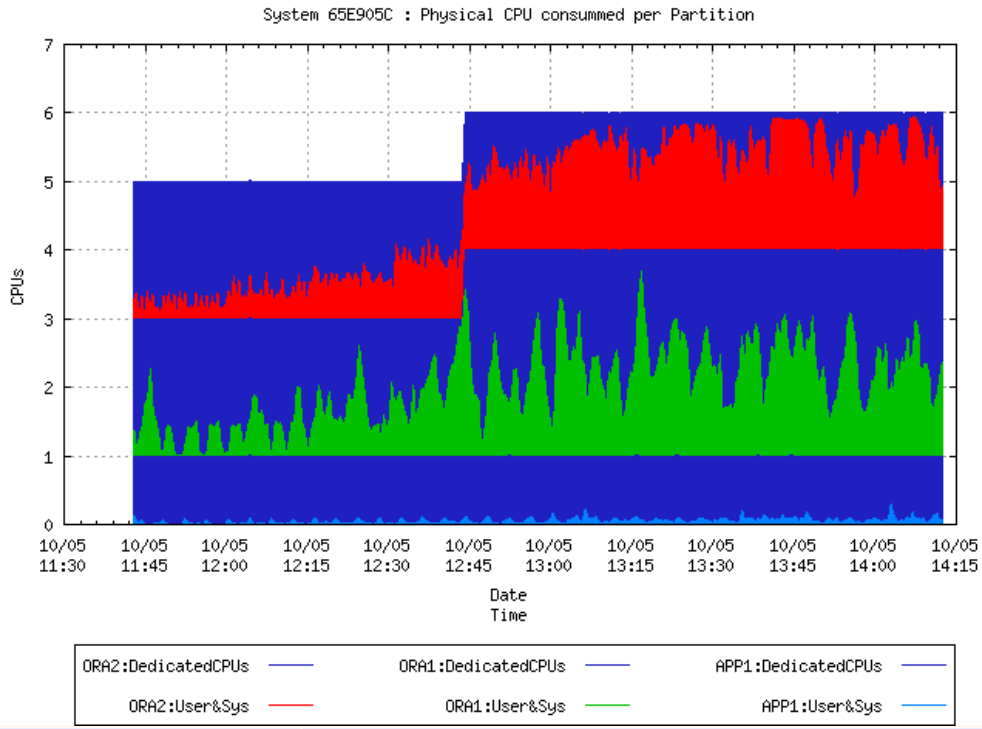

Shared : 6 shared processors, uncapped.

Transaction Response time (90%) : **8.636 sec**

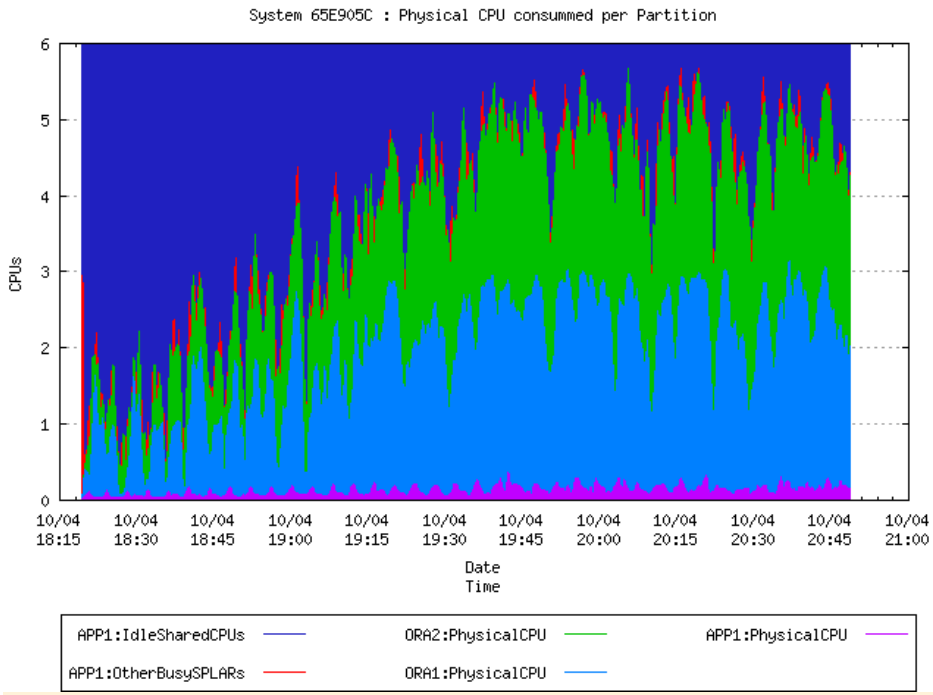

Oracle 9i benchmark: analysis

- Final configurations : the same
  - Dedicated : 6 CPUs Dedicated LPARs
  - Shared : 6 shared processors.
- Results :
  - CPU consumption ratio (C) =~ 90%
  - Processors configuration (P) =~ 100%
  - Better transaction response time for SPLPAR

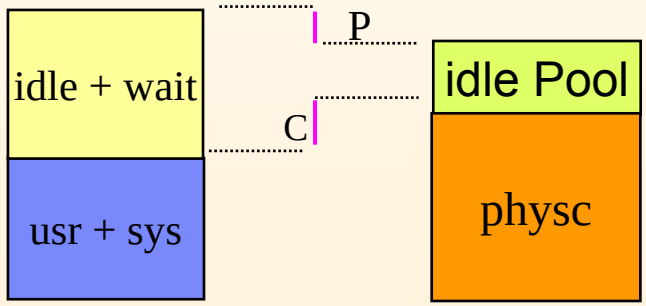

Oracle 9i benchmark: Analysis

- 2 Independent workloads : BO queries and Oracle Application Suite transactions
  - For the dedicated configuration, the Oracle Financial instance was 100%busy.
  - CPU consumptions Peeks of the two environments are asynchronous → better Processor Pool utilization

AMDOCS benchmark

- configuration :
  - 2 * P5 595
    - 64 Power5@1.9GHz
    - 256 GB RAM
    - 5 LPARs per system
  - 1 DS8300
  - AIX 5.3 ML3, VIOS 1.2
  - AMDOCS/Oracle 10g/TimesTen
- telco billing and rating: Messages passing application
- VIOS to share all network adapters:
  - EtherChannel (2 physical adapters)
  - 1SEA per VLAN

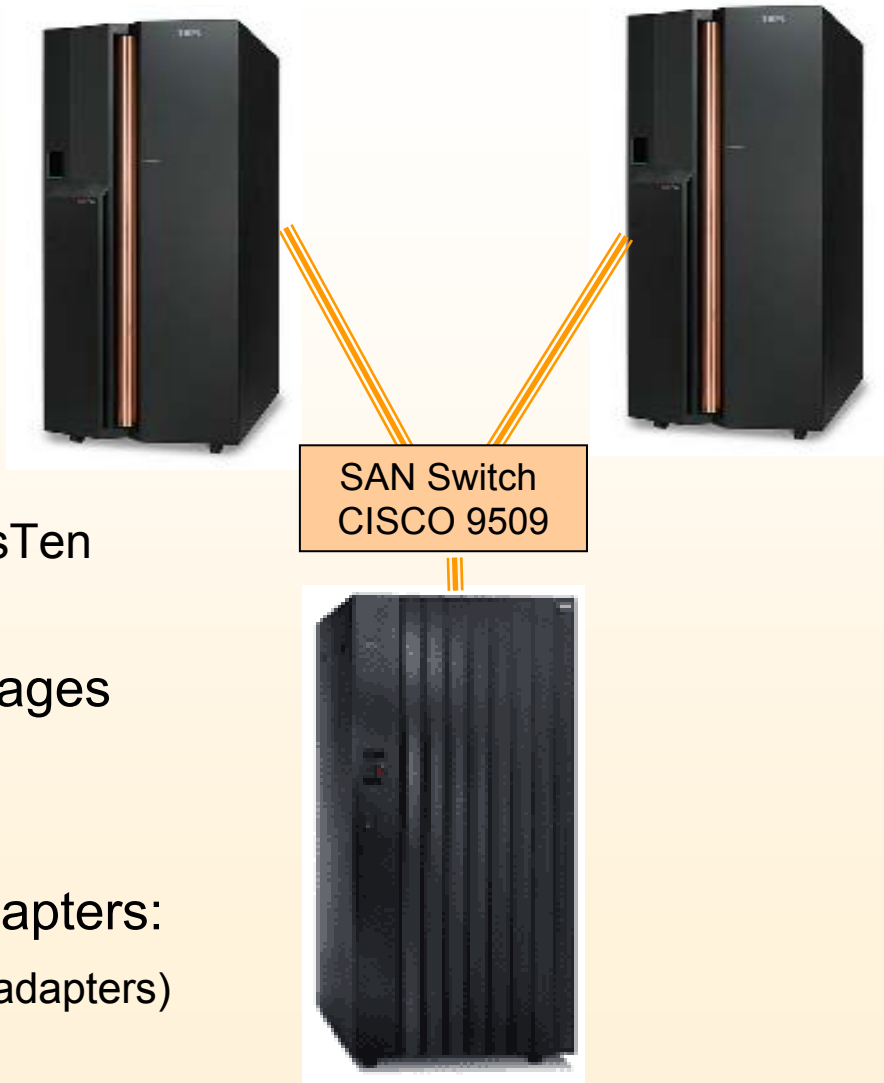

AMDOCS benchmark : network configuration

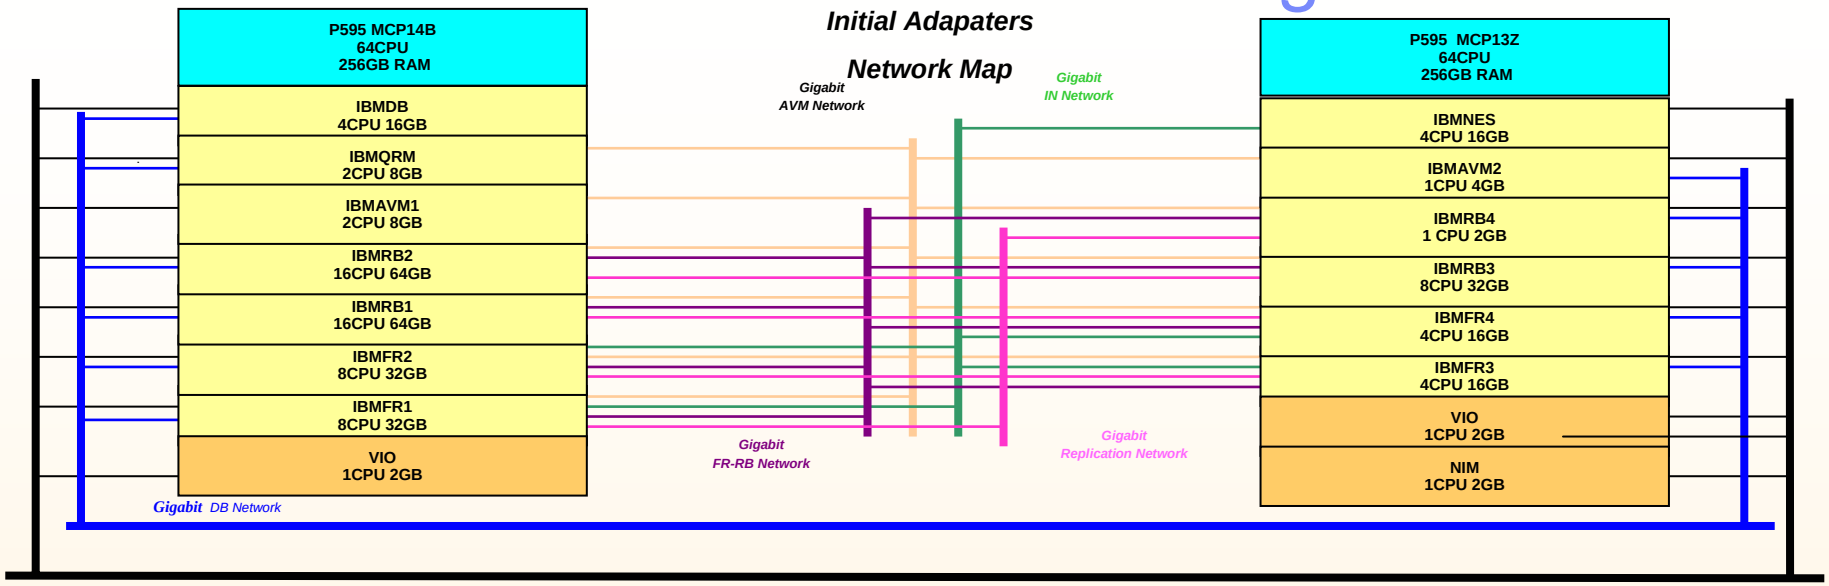

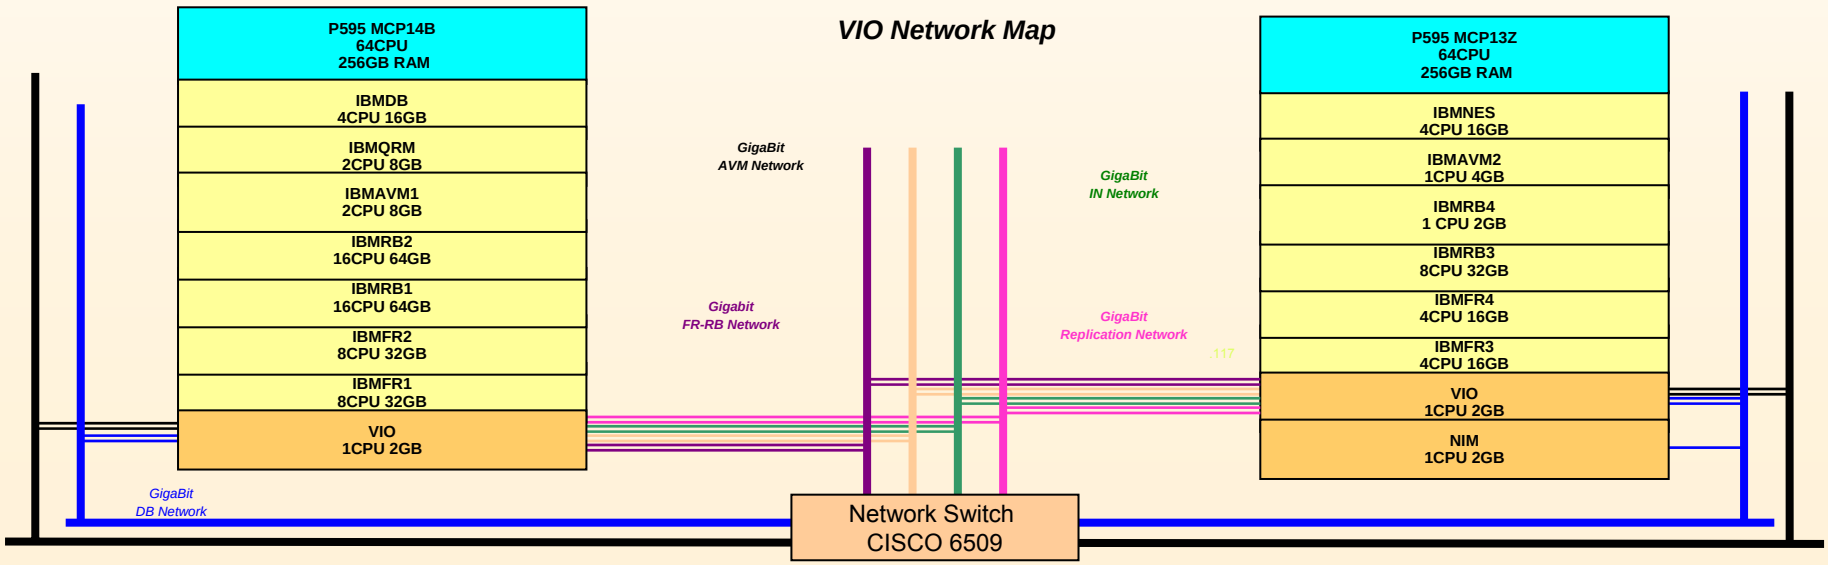

AMDOCS benchmark: results and analysis

- Results :
  - VIOS configuration :
    - 1 Cpu Dedicated
  - VIOS bandwidth →
  - VIOS CPU consumption :
    - ~70% CPU

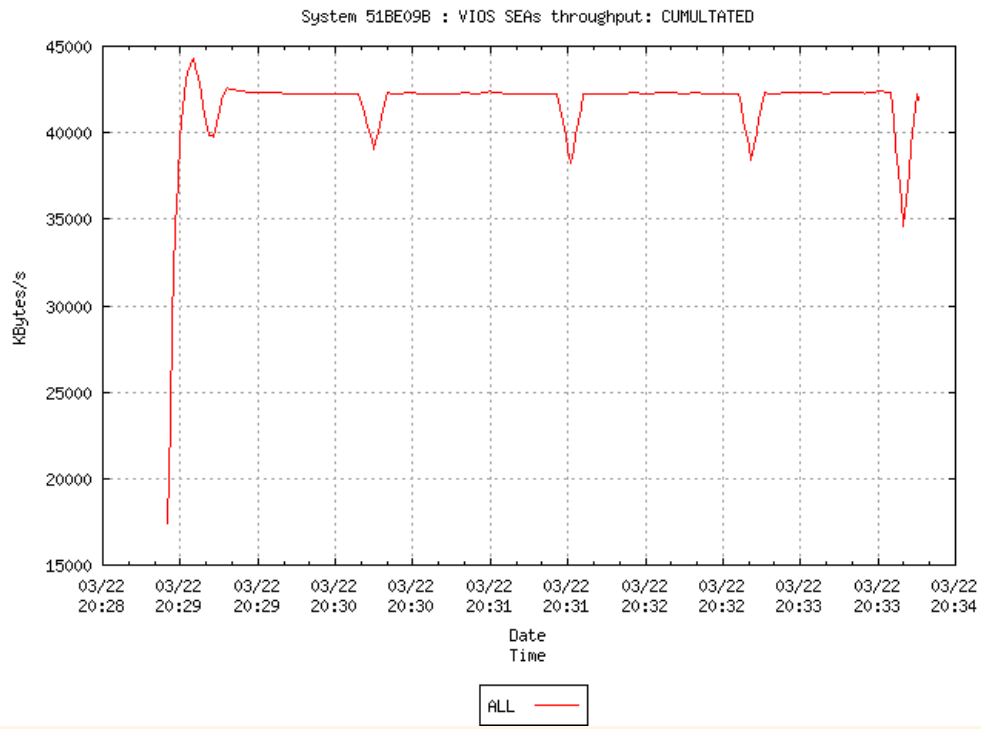

- Analysis :
  - SEA latency and throughput OK regarding AMDOCS expectations
  - Final tests reached the DS8k disks performance limits.

Virtualization Benchmarks : Lessons Learnt

- SPLPARs can definitely help to improve the application performance
  - Sizing results : SPLPAR resource cost =~ 15 to 20%.

- SPLPAR/DLPAR: Major impacts on performances
  - Independent workloads profile to optimize shared processor pool utilization.
  - Number of virtual processor in the Shared processors pool

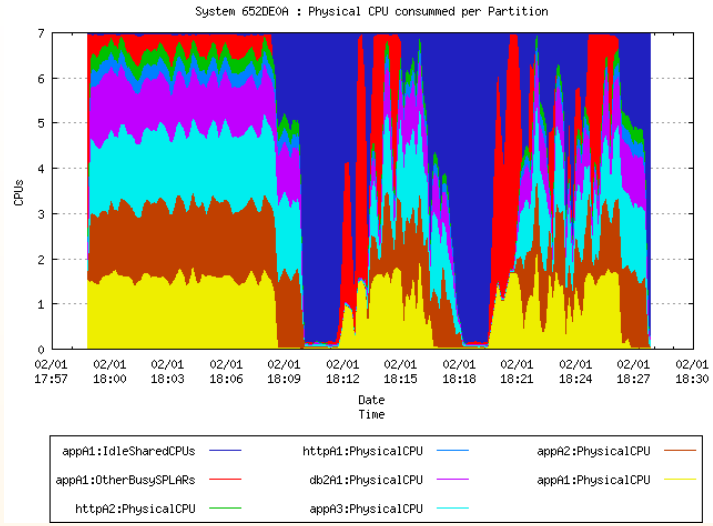

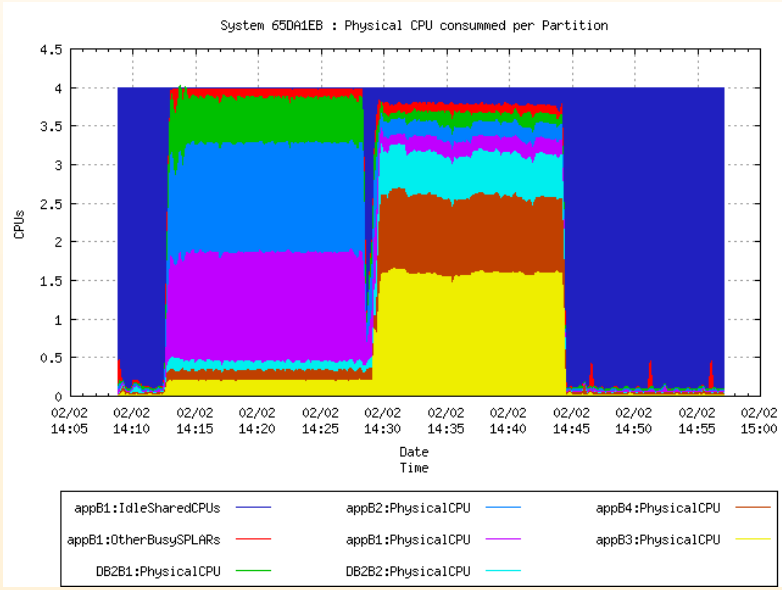

Agenda

- APV technology reminder
- Sizing and Benchmark results with APV solutions
- Customer and partner Production experiences
  - Consolidation example
  - Ephemeral environments/provisioning support
  - HA architecture, APV and CoD combinations

Agenda

- APV technology reminder
- Sizing and Benchmark results with APV solutions
- Customer and partner Production experiences
  - **Consolidation example**
  - Ephemeral environments/provisioning support
  - HA architecture, APV and CoD combinations

Bank : Business intelligence environment

- One Application → four environments
  - Production
  - Delivery
  - Pre-production
  - Development
- Requirements :
  - Price/performance
  - Flexibility and scalability

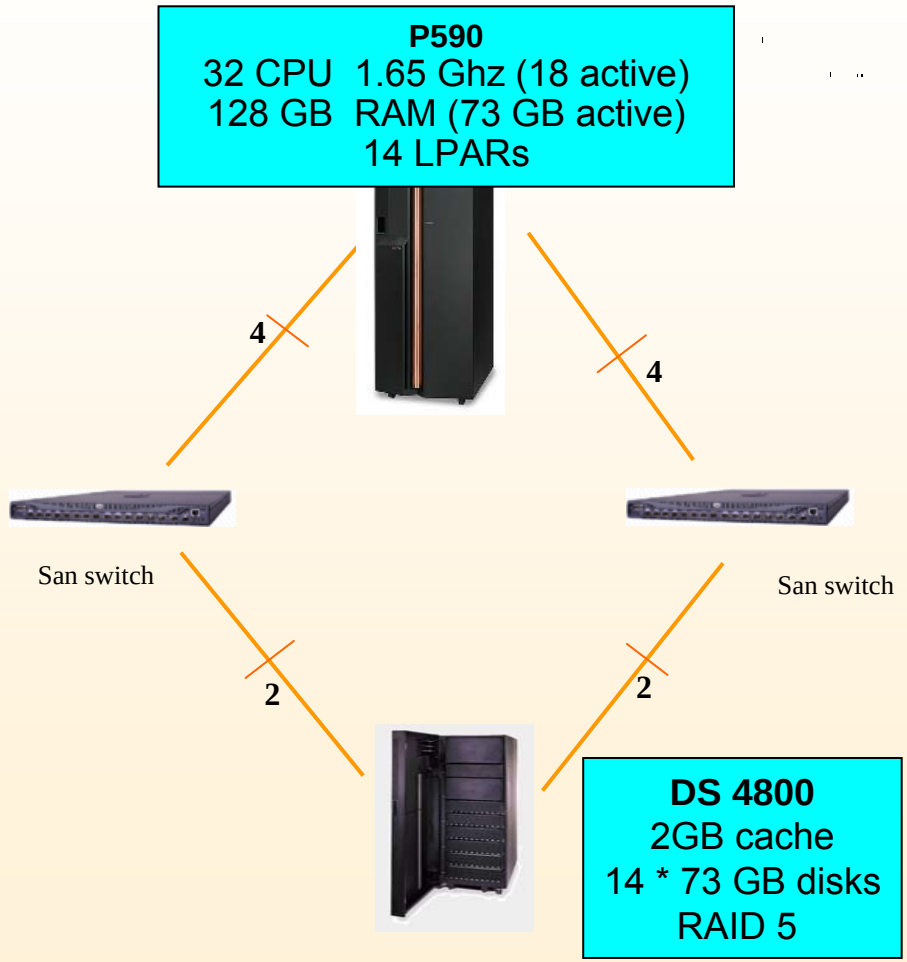

Bank : LPARs configuration

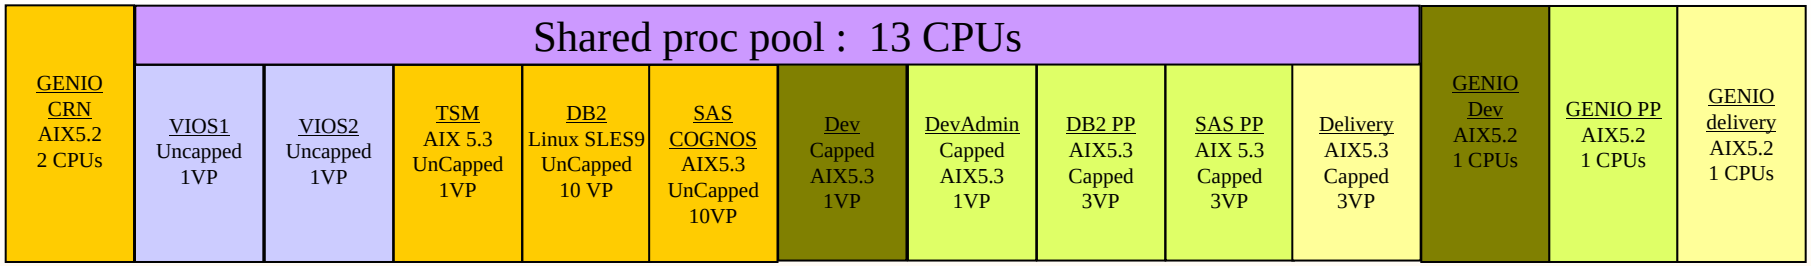

- 4 environments :
  - Production → Uncapped, enough VPs to consume Pool
  - Pre-Production → Capped
  - Delivery → Capped
  - Dev → Capped
- SPLPARs :
  - Dedicated LPAR required for Genio software (not released on AIX 5.3)
  - CE and VP are subject to changes, if the workload is required
  - Version : AIX 5.2 ML2 , VIOS 1.1

Bank : networks configuration

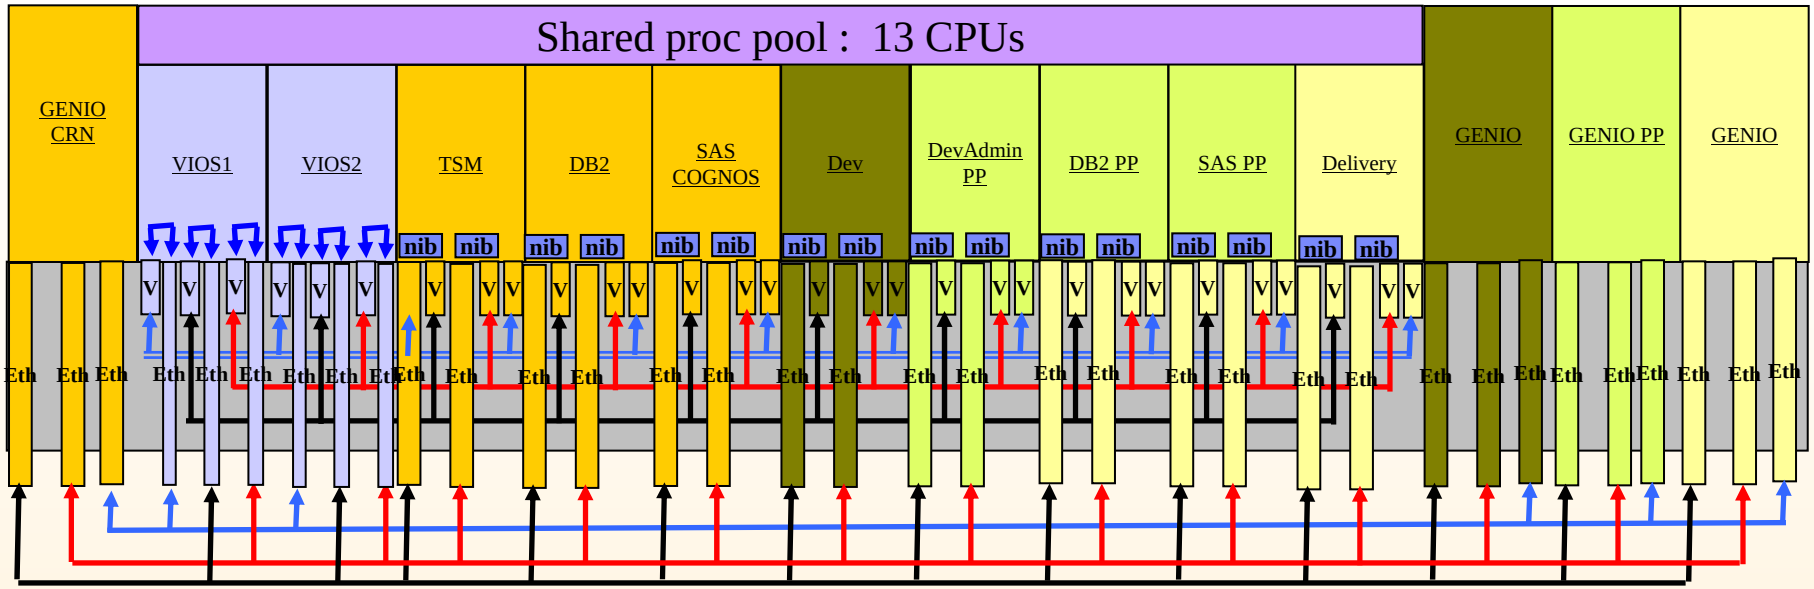

- 3 VLANs :
  - Admin (100Mb) All partitions, 2 Virtual VLAN (one per VIO) + client NIB,
  - Data(GigE)
  - Users(GigE)
- Physical Adapters : 17 dual Ethernet, 6 mono (VIOS)
- Use of AIX Network interface backup (nib), with physical adapter as primary, and virtual as secondary (same VLAN thanks to a VIOS SEA)
- No use of 802.1Q ethernet tagging.

Bank : disks configuration

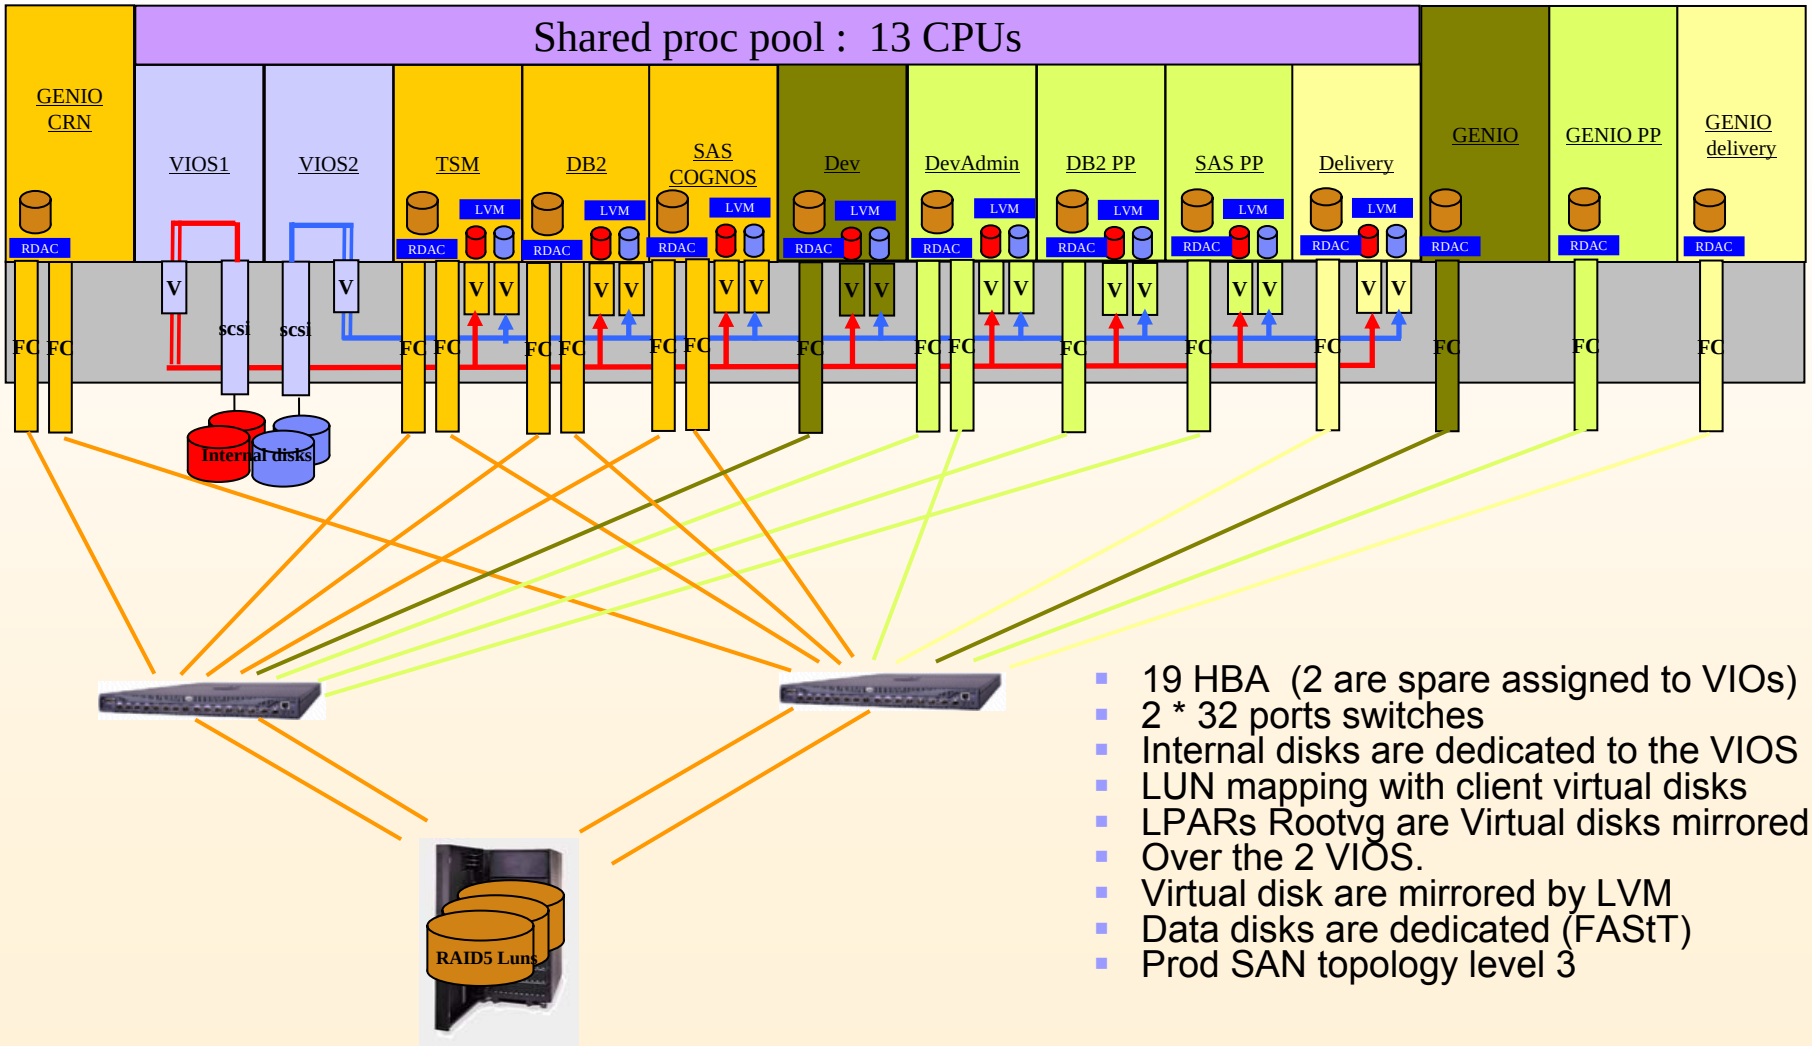

- 19 HBA (2 are spare assigned to VIOs)
- 2 * 32 ports switches
- Internal disks are dedicated to the VIOS
- LUN mapping with client virtual disks
- LPARs Rootvg are Virtual disks mirrored
- Over the 2 VIOS.
- Virtual disk are mirrored by LVM
- Data disks are dedicated (FAStT)
- Prod SAN topology level 3

Bank : Business intelligence

- Highlights
  - SPLPAR for AIX5.3, Dedicated LPAR for 5.2
  - Production LPARs are uncapped
  - Dev and Pre-Production are capped
  - VIOs for
    - Internal disks : LPARs Rootvg are all virtualized
    - Administration network SEAs
  - Data are stored on DS4800 LUNs and FC adapters are all dedicated to LPARs
  - Dedicated networks adapters for « users » networks
  - 1 SEA to support admin and data network : use of tcp/ip tagging
  - 14 inactivated Processors for potential “non disruptive” evolution
  - Network failover based on AIX NIB : dedicated Ethernet adapter as primary (for QoS) and a virtual Ethernet adapter for failover.

Agenda

- APV technology reminder
- Sizing and Benchmark results with APV solutions
- Customer and partner Production experiences
  - Consolidation example
  - **Ephemeral environments/provisioning support**
  - HA architecture, APV and CoD combinations

Software Vendor : media diffusion management

- Software Life cycle
  - Development
  - Testing
  - Integration
  - Customer situation Debug
- Quick software life cycles environments :
  - Creation, backup, deletion of environment
- Software :
  - Oracle 9i and 10g

Software Vendor : physical architecture

Phase 1 - since march 2005 :

- p550 4 CPUs 1.65GHz
- 13 SPLPARs in average
- Internal disks

Phase 2 – since end of 2005 :

- p570 8 CPUs 1.65GHz
- DS4800 disks + SAN switch

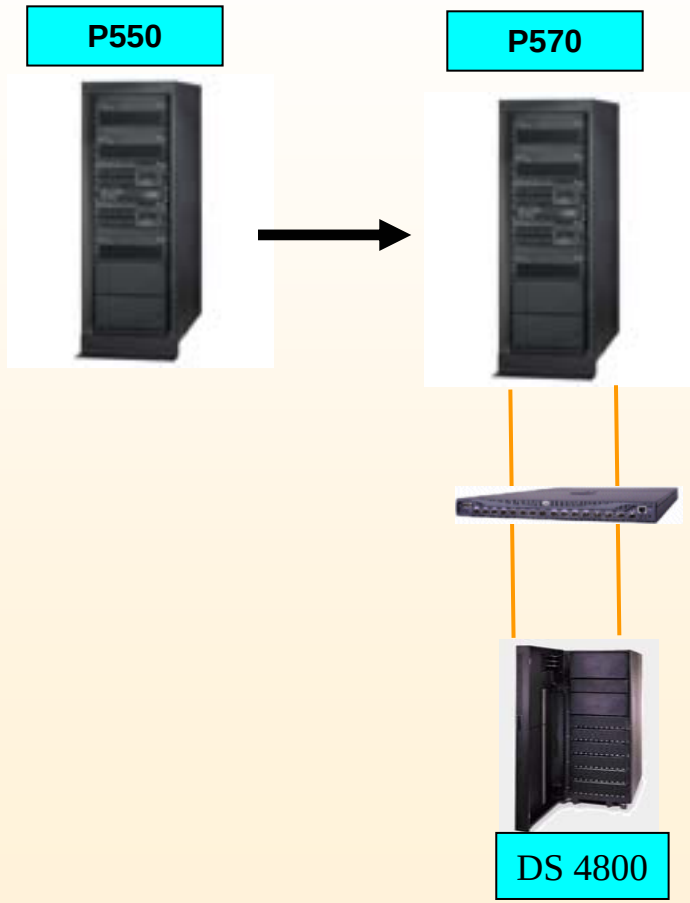

Software Vendor : configuration

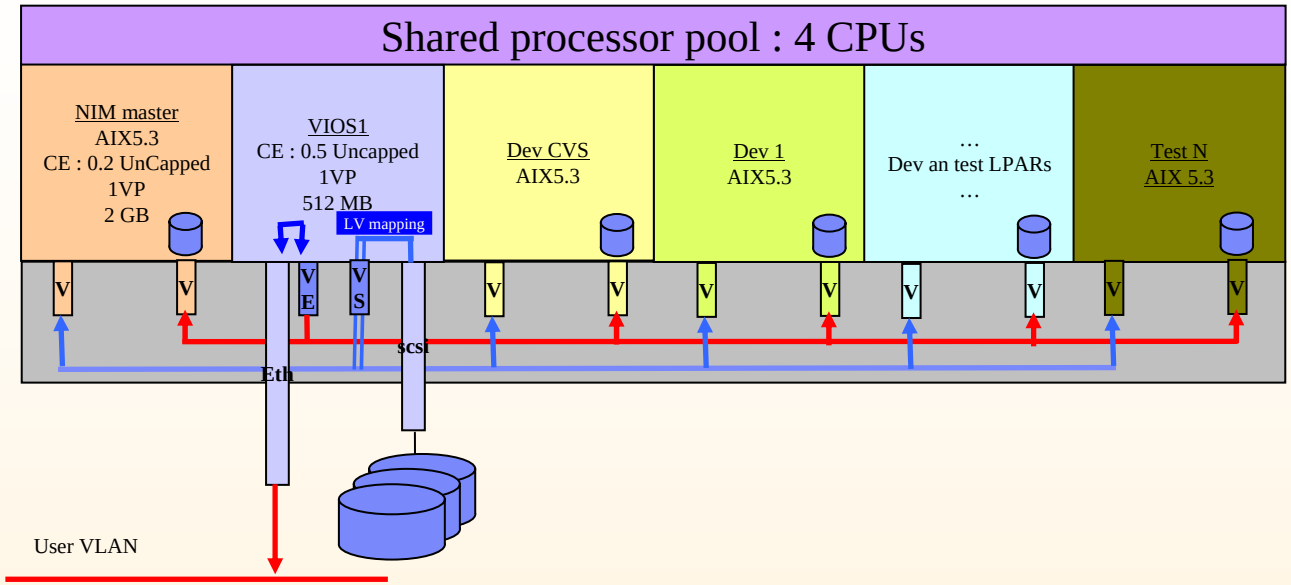

- All adapters are dedicated to the VIOS
- VIOS LVM slicing for virtual disks.
- All LPARs are uncapped micro-partitions
- Intensive use of NIM and alt_disk_install management to load, clone, backup and restore environment
- No VIOS protection nor HA networking yet
- Use of PLM in passive mode to monitor SPLPARs activity

Software Vendor example

- feedback
  - p5 technology very flexible (APV + NIM + Alt disk install)
  - Micro-partitions are ideal for Development environment
    - 4 way p570 capacity is actually enough for ~15 developers and ~10 testers.

pLite

- Plite: Light system p5 customer benchmark
  - UIF Session “Plite: Transforming benchmarking with VE”
- Based on ODINA (On Demand Infrastructure Automation):
  - Reservation (based on hard-ware provisioning)
  - Automatic platform Set-up (as far as possible)
  - I, p, z and x platform support : p5 system fully supported
    - Implemented with Tivoli Provisioning Manager (TPM), CSM, NIM, HMC scripting, VIO 1.2
- Several other usages :
  - Demo and Light Proof of Concept environment
  - Service provided since September 2005
    - We run 2 pLite benchmarks in parallel in average

Agenda

- APV technology reminder
- Sizing and Benchmark results with APV solutions
- Customer and partner Production experiences
  - Consolidation example
  - Ephemeral environments/provisioning support
  - **HA architecture, APV and CoD combinations**

Media provider

Infrastructure evolution for a application for subscribers management and its associated decisional application.

- Subscribers management application :
  - Applications SIGMA,
  - PPV (Pay Per View),
  - Interactive Applications (Satellite and ADSL)
  - Oracle 9i
- Decisional analysis
  - DW DB: Oracle 9i
  - data extraction based DataStage
  - User requester : ESSBASE, SAS and BO
  - Oracle BD or data warehouse
  - Cubes ESSbase

Media provider: description of the new solution

- Description
  - 2 physical servers for all environments with CoD On/Off
  - CoD On/Off activated for heavy but unusual or HACMP failover.
  - HACMP cross failover for 2 environments (symmetry)
  - HACMP activates On/Off when taking over through HMC CLI scripts
  - Internal disks managed by VIOS for rootvgs LPARs
- Advantage :
  - Server configuration reduced (CPU resources for taking over applications are activated and charged only when needed).
- Constraints
  - With the version of HACMP, backup NIB had to be physical Ethernet adapters.

Media provider : physical infrastructure

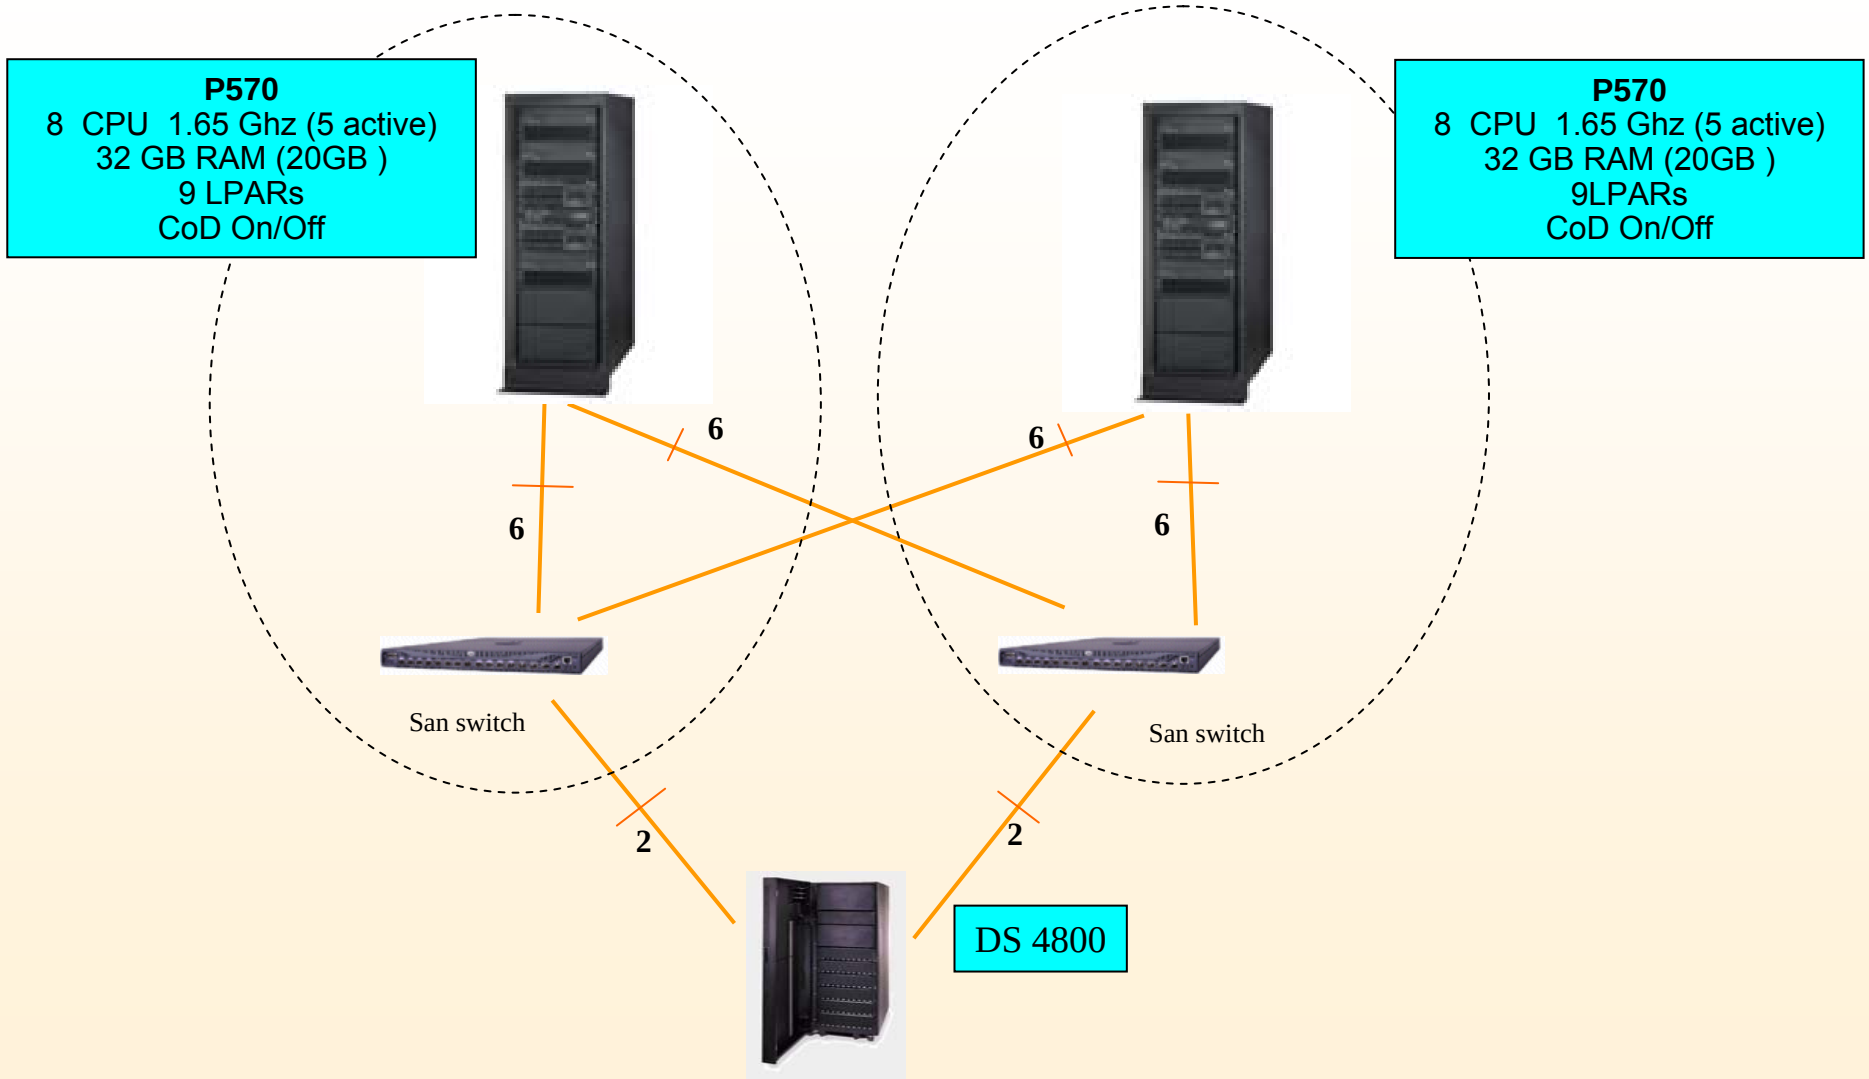

Media provider: SPLPAR configuration

P570
8 CPU 1.65 Ghz (5 active)
32 GB RAM (20GB)

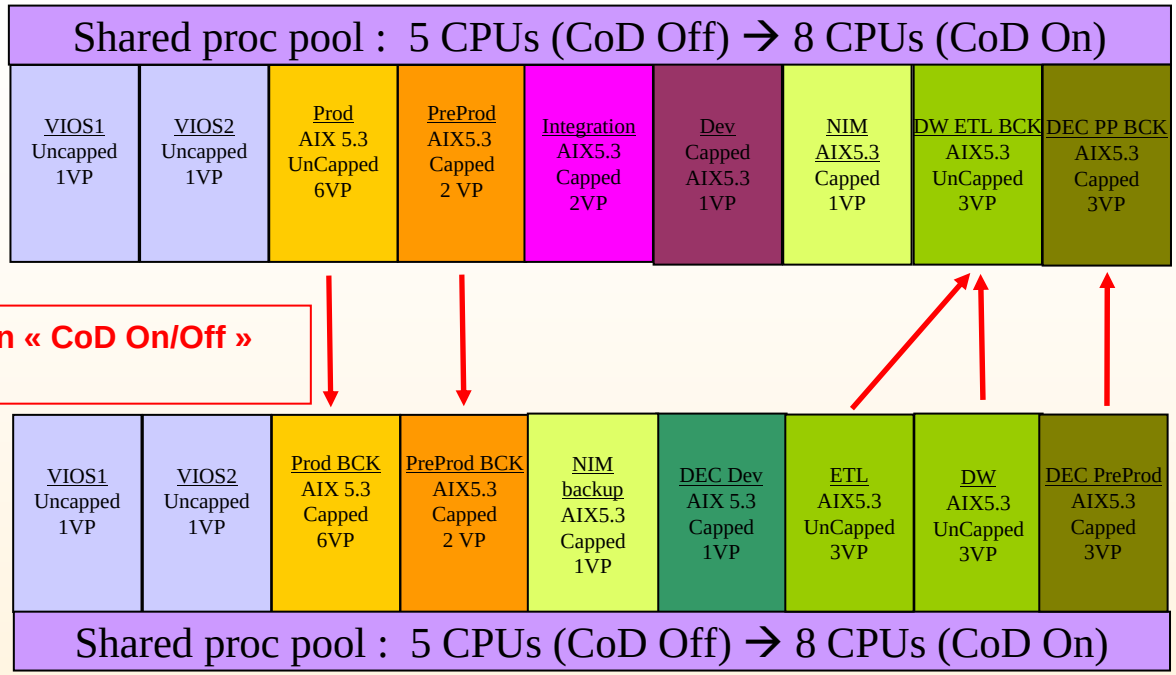

P570
8 CPU 1.65 Ghz (5 active)
32 GB RAM (20GB)

- **2 applications :**
  - **Subscribers management**
    - **Production + BCK** → **Uncapped, enough VPs to consume Pool**
    - **Pre-Production** → **Capped**
    - **Dev** → **Capped**
    - **Integration** → **Capped**
  - **Decisional analysis**
    - **Production + BCK** → **Uncapped, enough VPs to consume Pool**
    - **PreProd** → **Capped**
    - **Dev** → **Capped**

Media provider : network configuration

- For SPLPARs with HACMP:
  - 2 physical Ethernet cards for each user network using HACMP Stand-by.
- For SPLPARs without HACMP:
  - 1 Virtual Ethernet adaptor linked to a VIOS SEA
  - The public Admin HMC & TPS Backup network is virtualized.
- For each VIOS, 2 physical Ethernet Cards

Media provider: Analysis

Highlights :

- 2 physical servers for all environments with CoD On/Off
- CoD On/Off activated for heavy but unusual workloads or HACMP failover.
- HACMP activates On/Off when taking over through HMC CLI scripts
- Internal disks managed by VIOS for rootvgs LPARs

Main drawbacks

- VIOS monitoring
- HACMP backup NIB have to be physical

Pharmaceutical laboratories

- New DB infrastructure : consolidation and new application deployment
  - consolidation of Oracle RDBMS coming from different geographical regions
  - Deployment of a new data warehouse application
    - Support of 4 environment per application :
    - Development, test, production and backup
  - Evolution : scale from 1 to 2
  - Software stack :
    - Oracle 10g
    - Business Objects

Pharmaceutical laboratories: solution description

- Baselines :
  - Use of On/Off CoD from the Production environment
  - Dissymmetric servers configuration
  - Use of Reserve CoD from the HACMP standby server :
    - Permanent CPUs number is half of Primary server CPUs number but it can grow to the same size when taking over.
  - Use of NIB with a virtual Ethernet adapter as the backup interface

Physical infrastructure

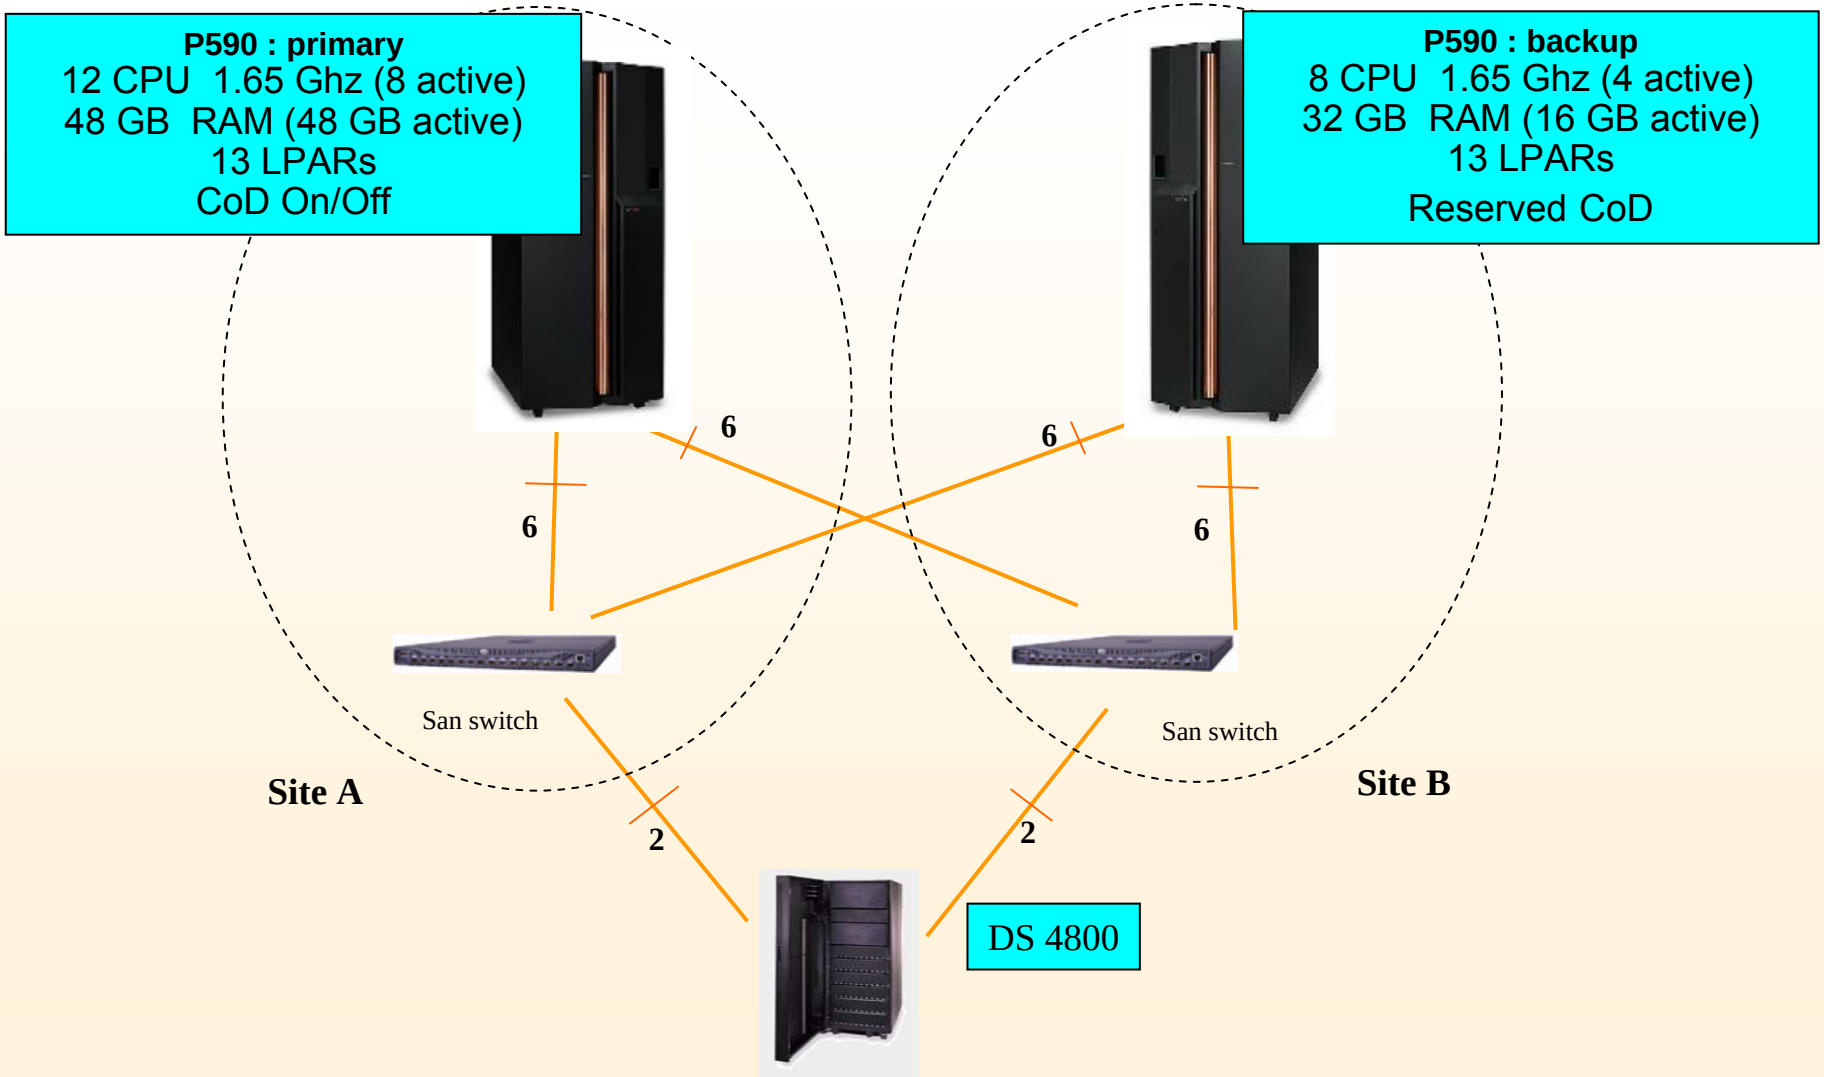

SPLPARs configuration

P590
- 12 CPU 1.65Ghz -8 active
- 48 GB RAM -48 GB active
- 13 LPARs
- CoD On/Off

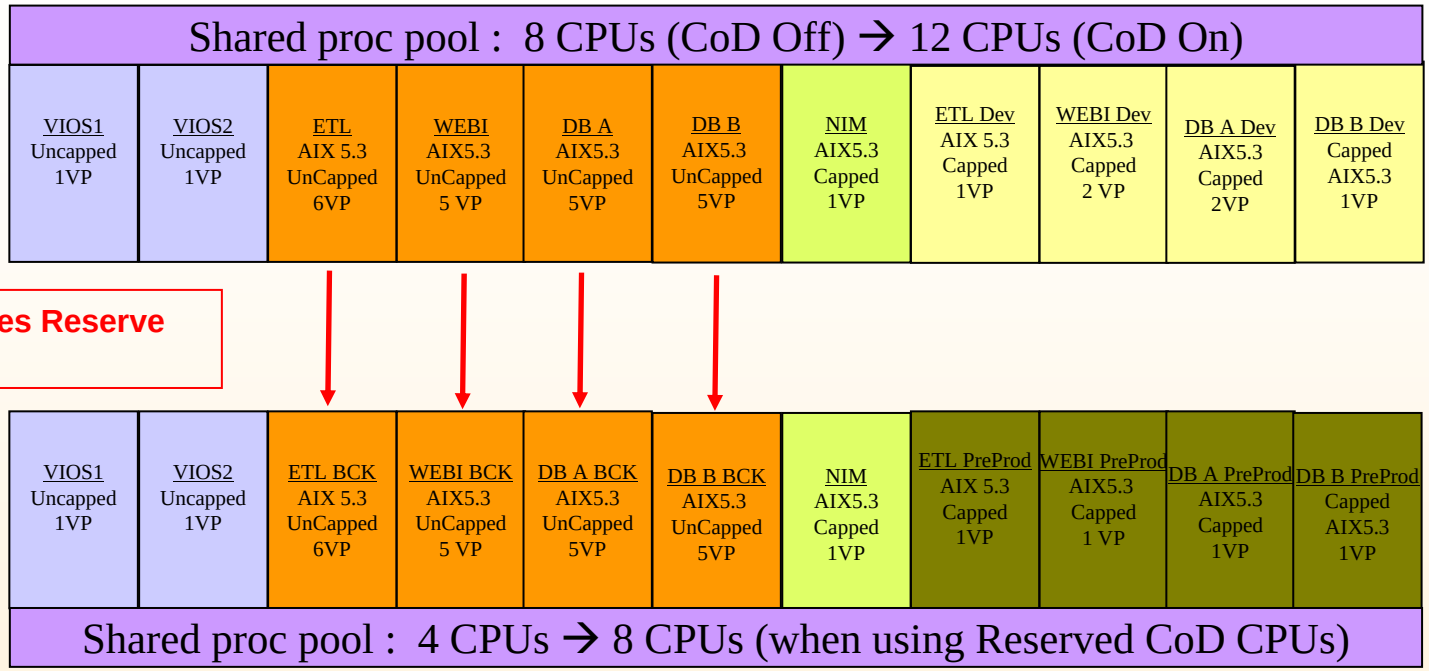

P590
- 8 CPU 1.65 Ghz (4 active)
- 32 GB RAM (16 GB active)
- 13 LPARs
- Reserved CoD

- 4 environments :
  - Production → Uncapped, enough VPs to consume the Pool
  - NIM → Capped
  - Dev → Capped
  - Pre Production → Capped , to prevent Reserve CoD activation

Pharmaceutical laboratories : Analysis

Highlights:

- **Variable CPU** capacity for the **backup** server :
  - permanent configuration is ½ the production configuration
  - It can grow to the same size with Reserve CoD in failover mode. When HACMP failover on the backup server and the application workload increases again, the Reserved CoD CPUs are used automatically.
  - To prevent Reserve CoD CPUs to be used in non failover mode, the other environment are **capped** or have a reduced number of VPs
- Primary server : **CoD On/Off** activated for heavy but **unusual workloads**.
- Internal disks managed by VIOS for the rootvg disks of client LPARs
- physical HACMP backup network interface.

Main drawbacks:

- VIOS monitoring

Trends collected through these experiences

- Capped/Uncapped mode:
  - Production SPLPAR Uncapped
  - VIOS SPLPARs Capped
- Virtual/Dedicated Disks :
  - Rootvg disks: Virtual, whole LUN
  - Data disks: Dedicated
- Virtual/dedicated Network adapters
  - Administration network : Virtualized
  - Production : use of NIB with dedicated NIC as primary and Virtual as backup
- VIOS : doubled
- HACMP.
  - Use of CoD On/Off or Reserve CoD
- CoD :
  - On/Off : for heavy and infrequent workload or HACMP
  - Reserve CoD : combination with HACMP (but other active SPLPARs has to be Capped)

Questions …

- **Thanks a lot**