vPars and HP-UX Adaptive Infrastructure

Henrik Elmgren
Solution Architect
HP Competitive Sales/Pre-Sales

Alan Hymes
Solution Architect
HP Competitive Sales/Pre-Sales

HP WORLD 2003
Solutions and Technology Conference & Expo

Business needs demand a new model of computing

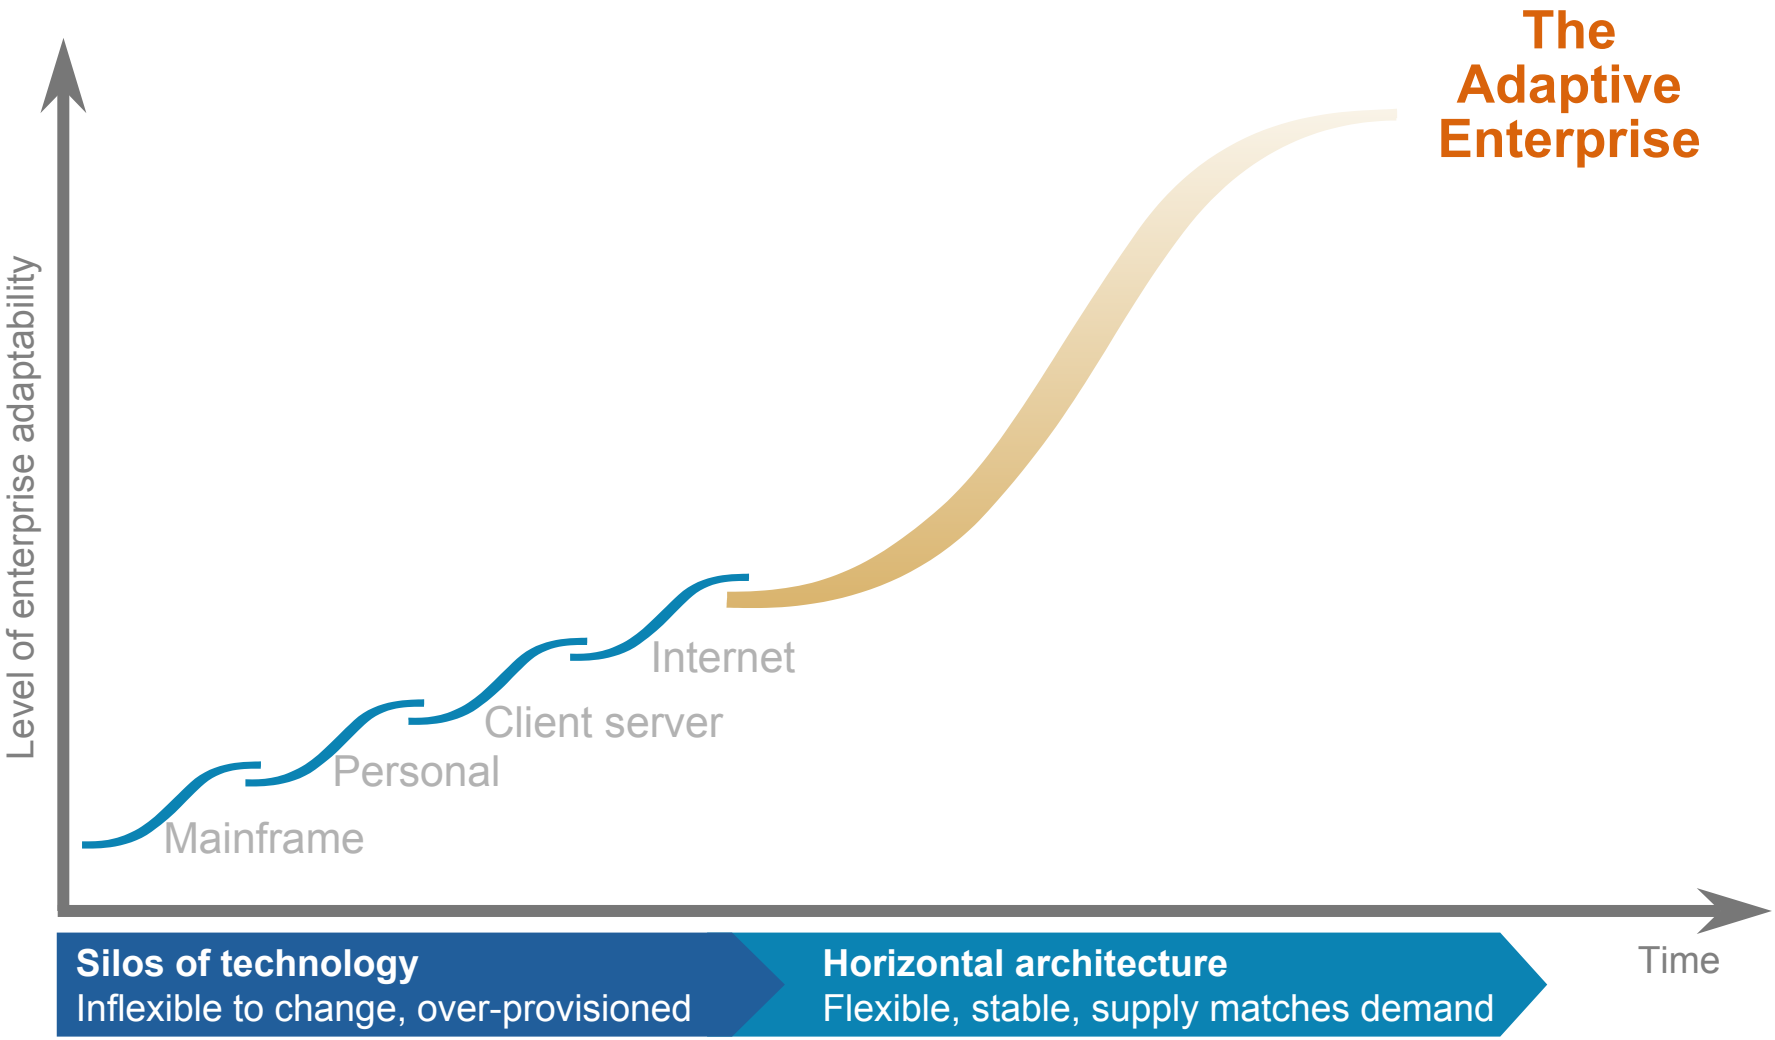

Virtualization Enables the Adaptive Enterprise

Ideal computing

- Pay for what you use
- Share virtual resources

Result: IT shifts focus from boxes to services

HP Virtual Server Environment and HP UDC Evolution

On the road to utility computing

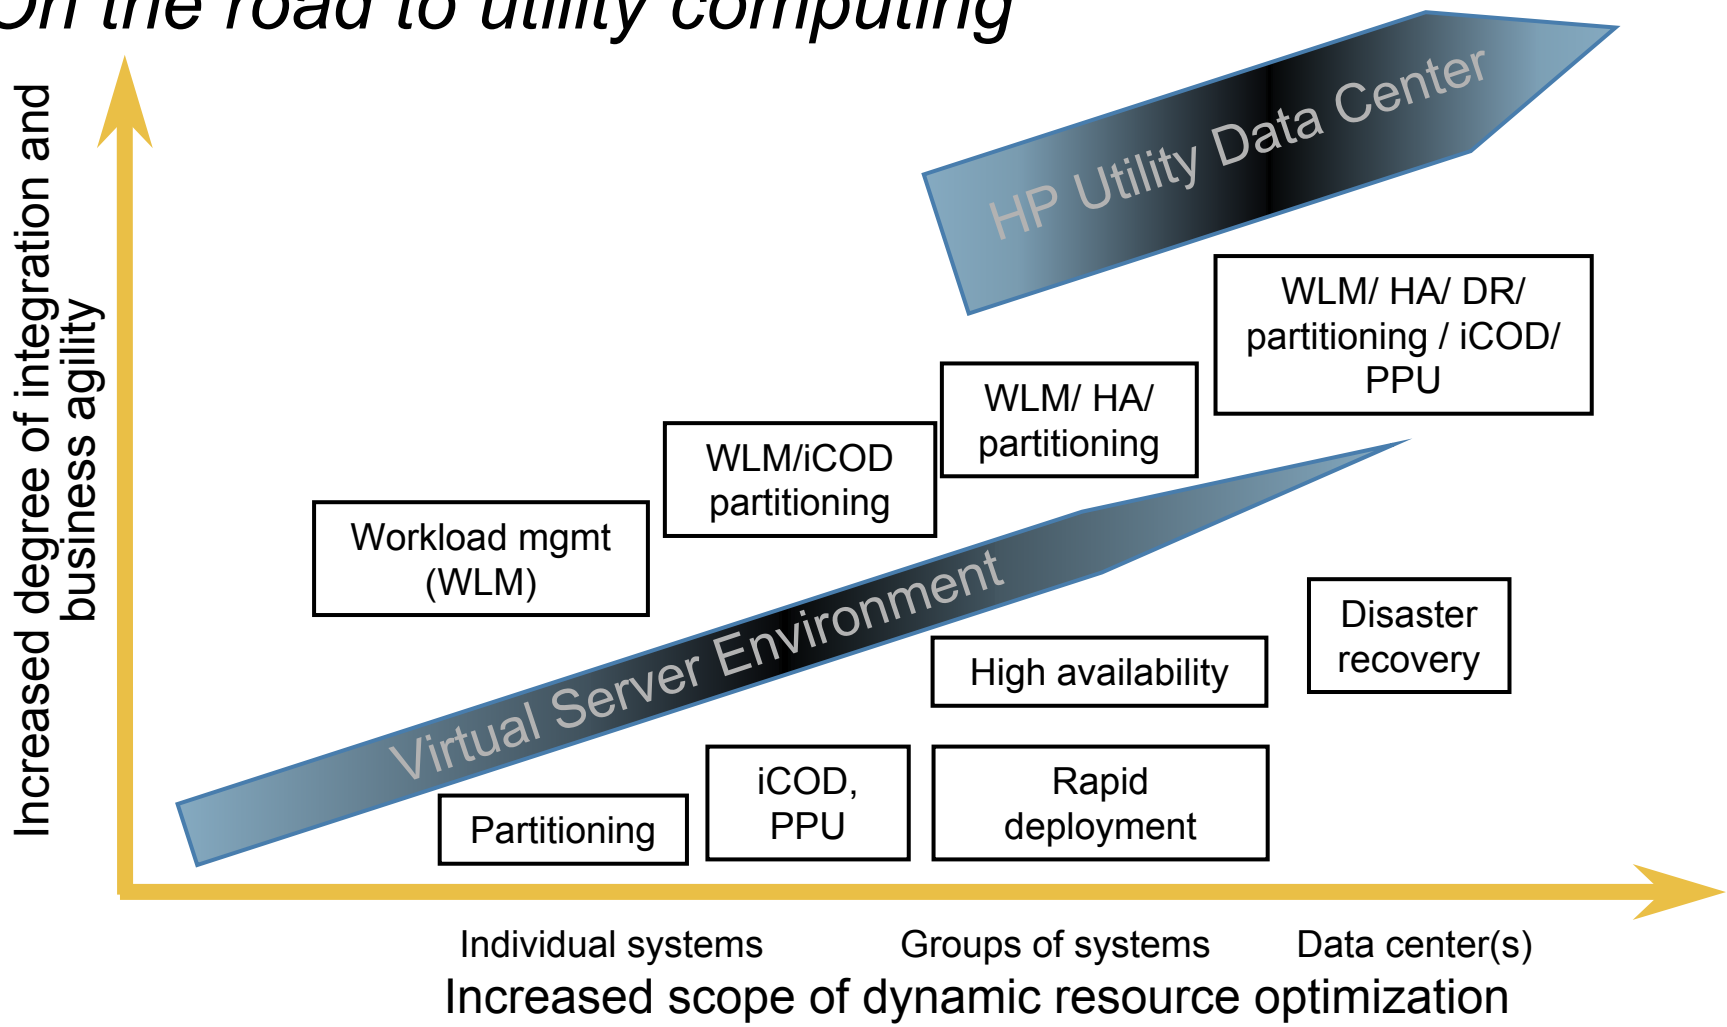

HP Virtual Server Environment

Based on the only UNIX goal-based policy engine

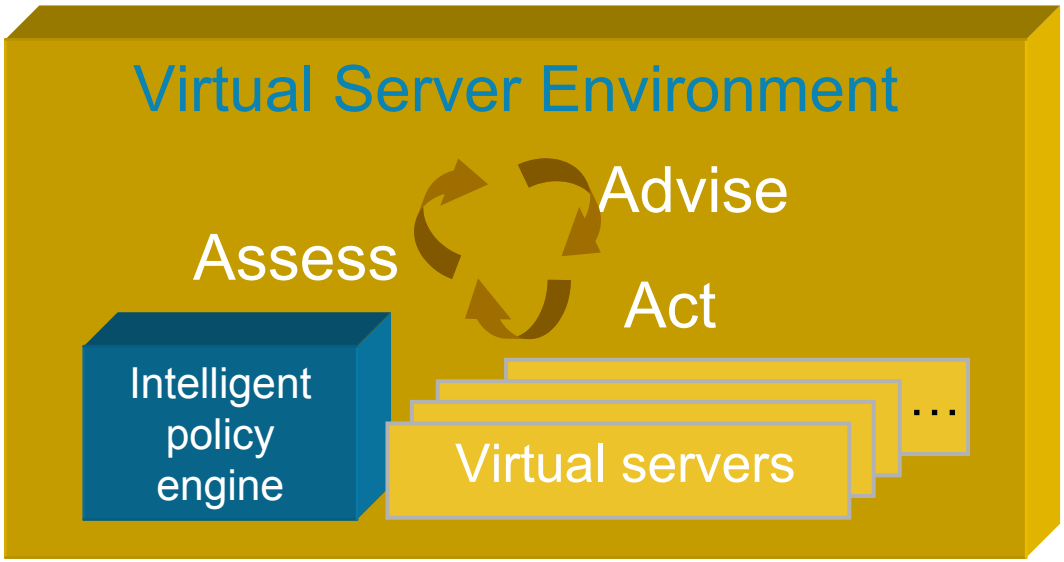

Expands and shrinks virtual servers in real time based on business priorities

- Better RoIT through optimized resource utilization
- Increased business agility through the capability to allocate resources on the fly
- Highest Quality of Service through continuous real time assessment, advice, and action

HP Virtual Server goals

Customer issues today!		***How HP-UX server virtualization addresses***
Pressure to offer service level guarantee at reasonable costs	→	Meets service level agreements with best return-on-investment
Under utilization of servers	→	Improves utilization of server compute power up to 80/90%
Address high fluctuation of Web and App traffic	→	Provides fast and dynamic implementation for changing requirements
Flexibility with privacy and high availability	→	Provides “right” level of application isolation with uptime

HP Virtual Server Environment for HP-UX:

Intelligent orchestration of virtualized server resources

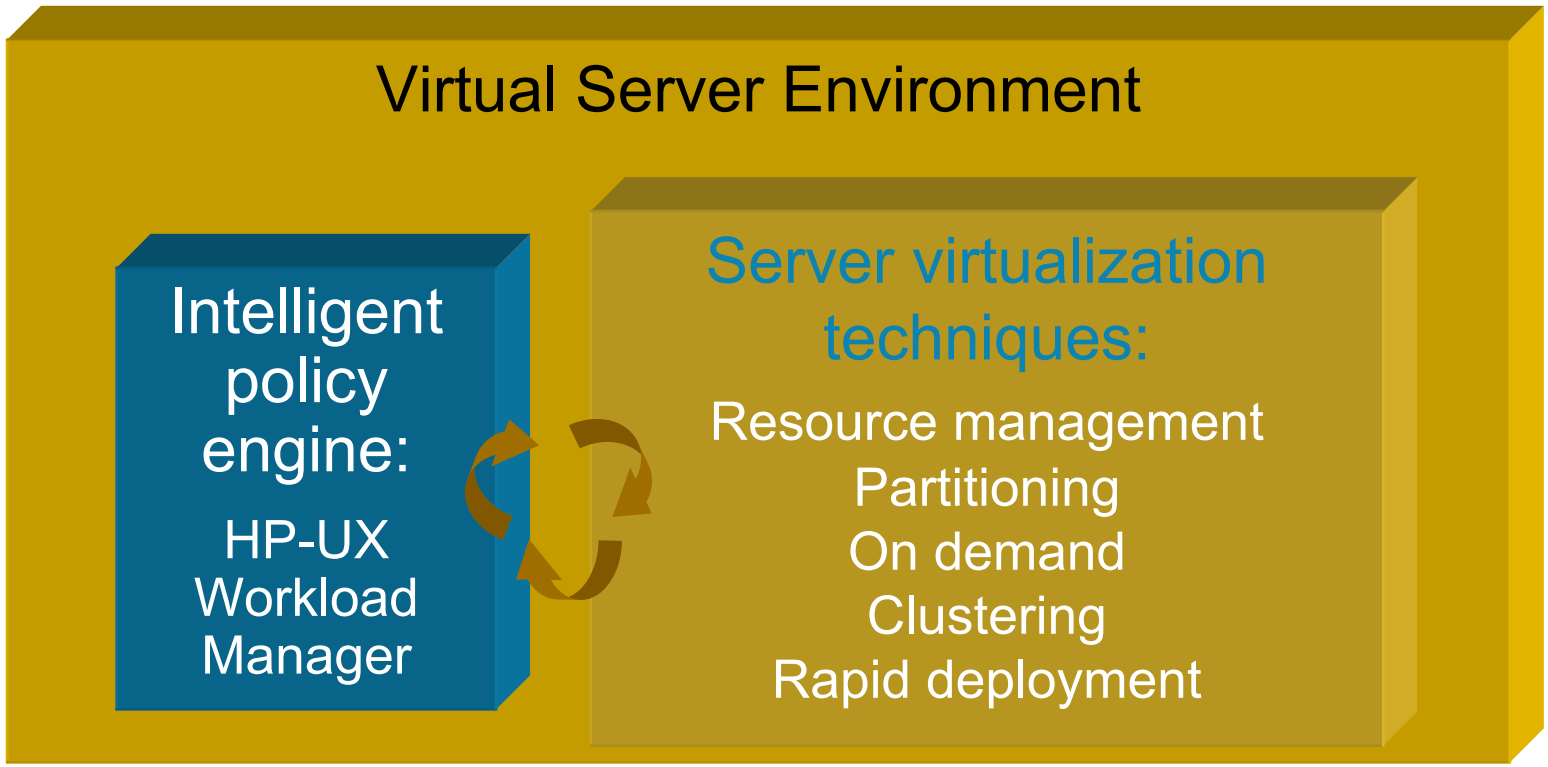

- Increased agility through tight integration of goal-based policy engine with virtualization techniques
- Reduced complexity through integration with cost-effective multi-system management tools

HP Partitioning Continuum for HP-UX

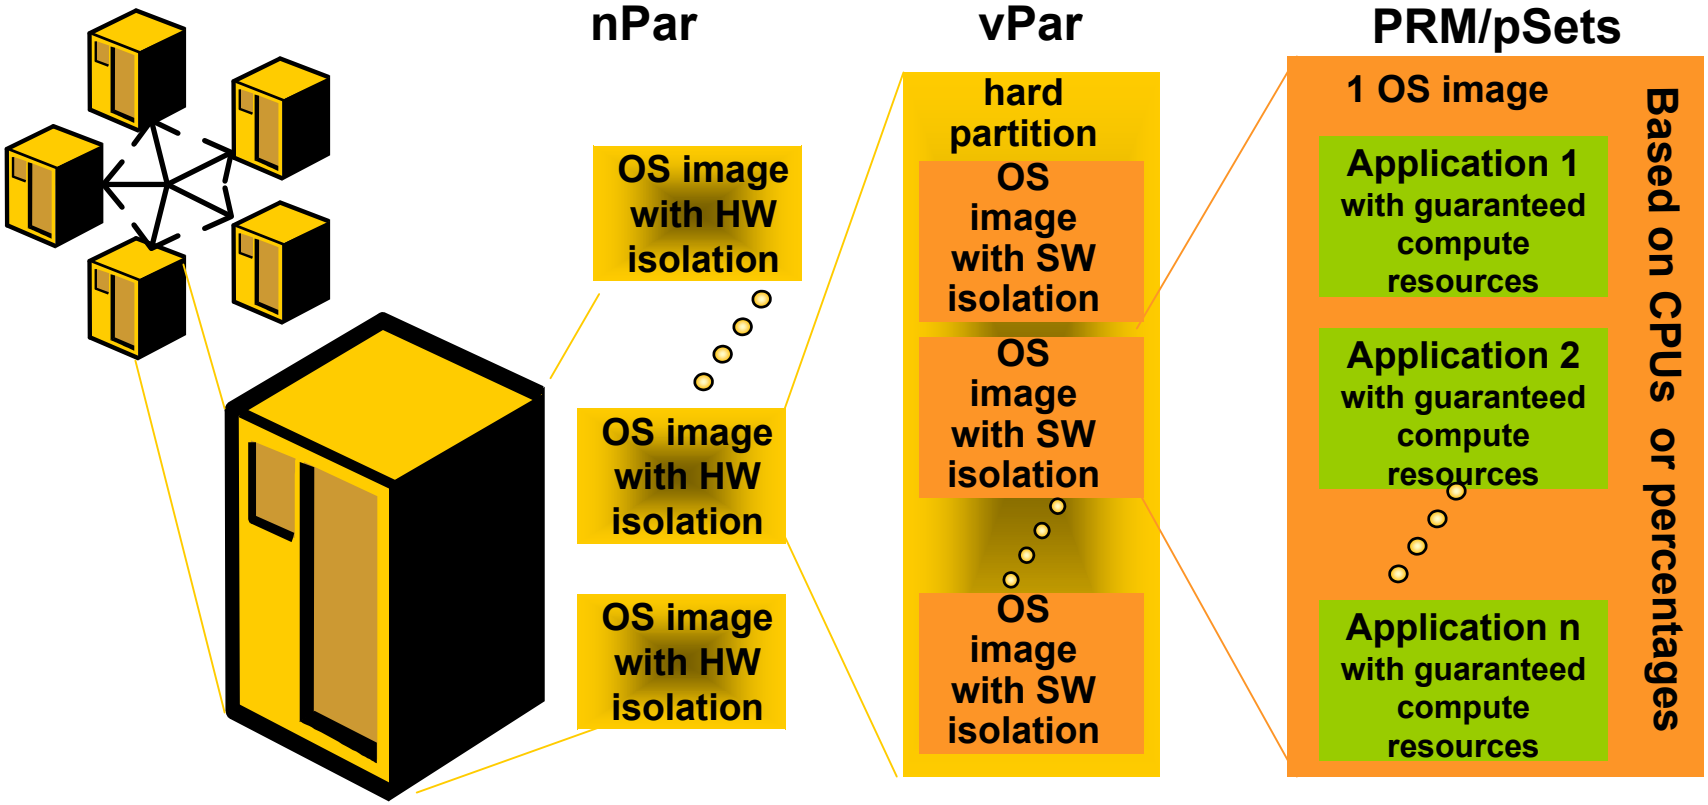

HP-UX Workload Manager

Isolation
Highest degree of separation

Flexibility
Highest degree of dynamic capabilities

nPartitions (hard partitions)

Multiple O/S instances per node with hardware isolation

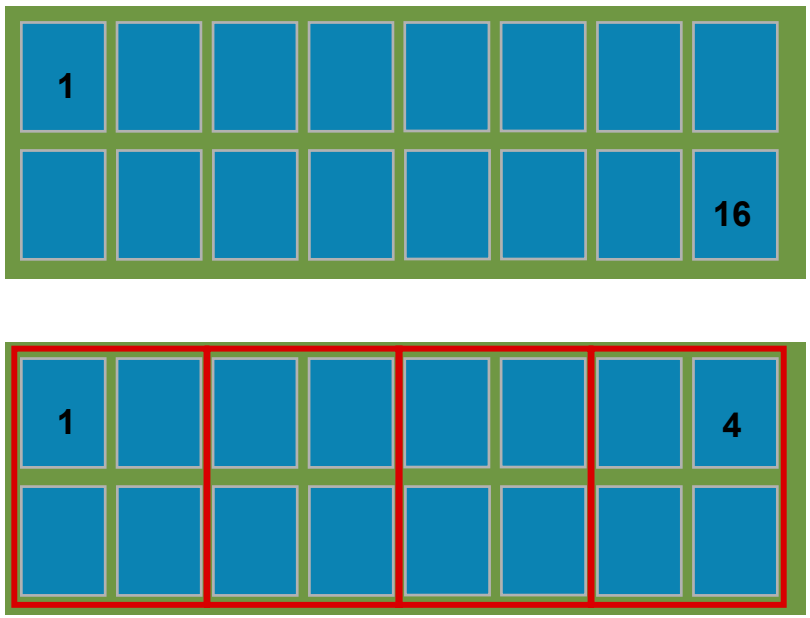

- Increased system utilization
  - Partitioning a server increases the utilization level. A Superdome can have up to 16 nPartitions
- Increased Flexibility: Multi OS
  - Multi OS support: HP-UX, Linux (*), Windows (*)
  - Multi OS version and patch level support
- Increased Uptime
  - Hardware (electrical) and software isolation across nPartitions
  - Serviceguard support (within a Server or to another HP 9000 server)
- Available on Superdome, rp8400 and rp7410

rp7410
nPartitions architecture

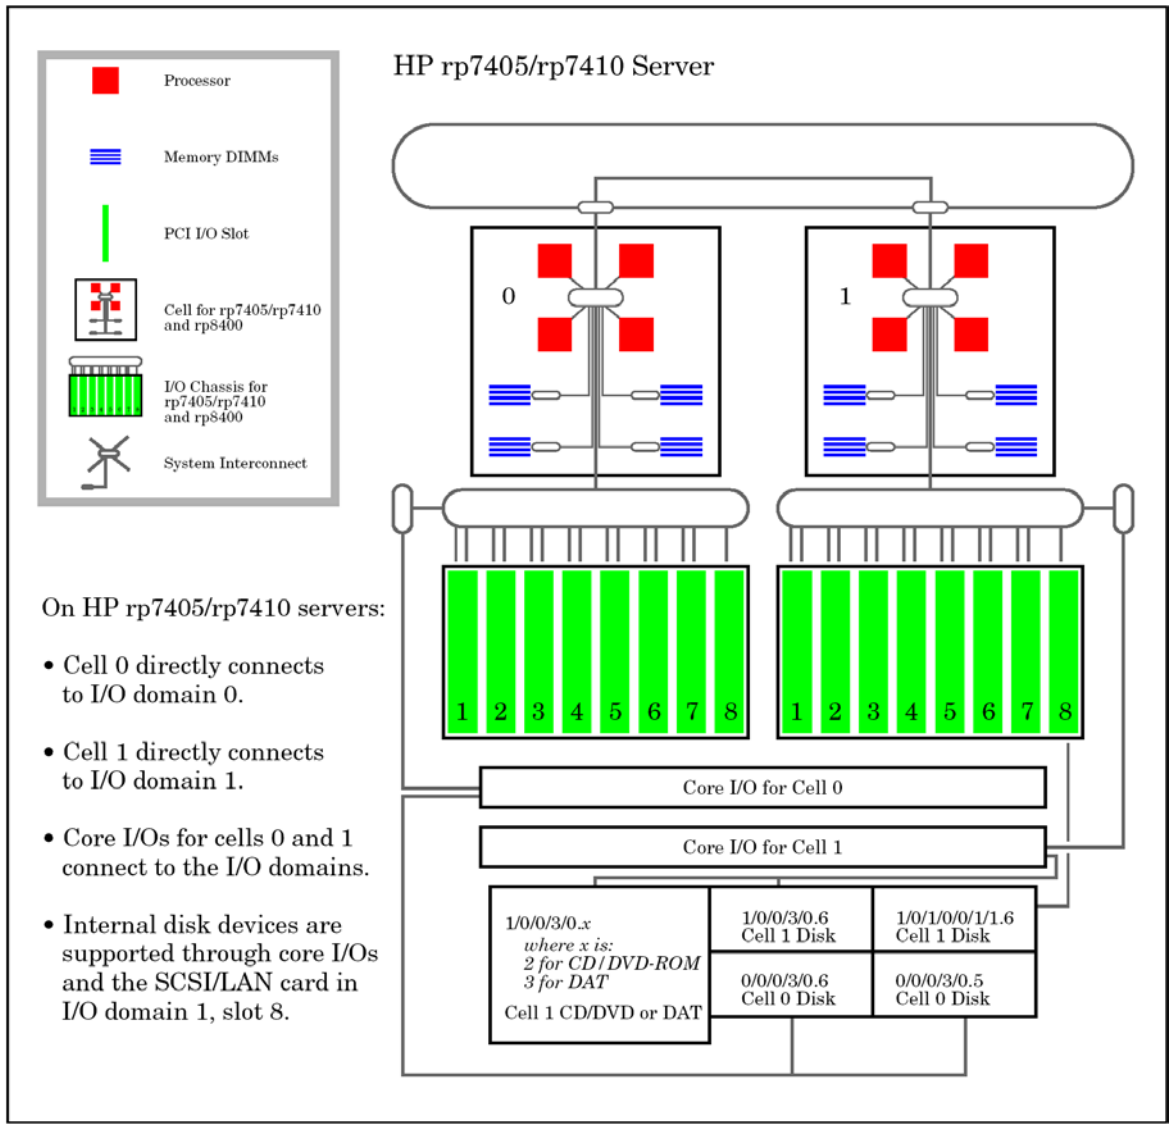

rp8400
nPartitions architecture

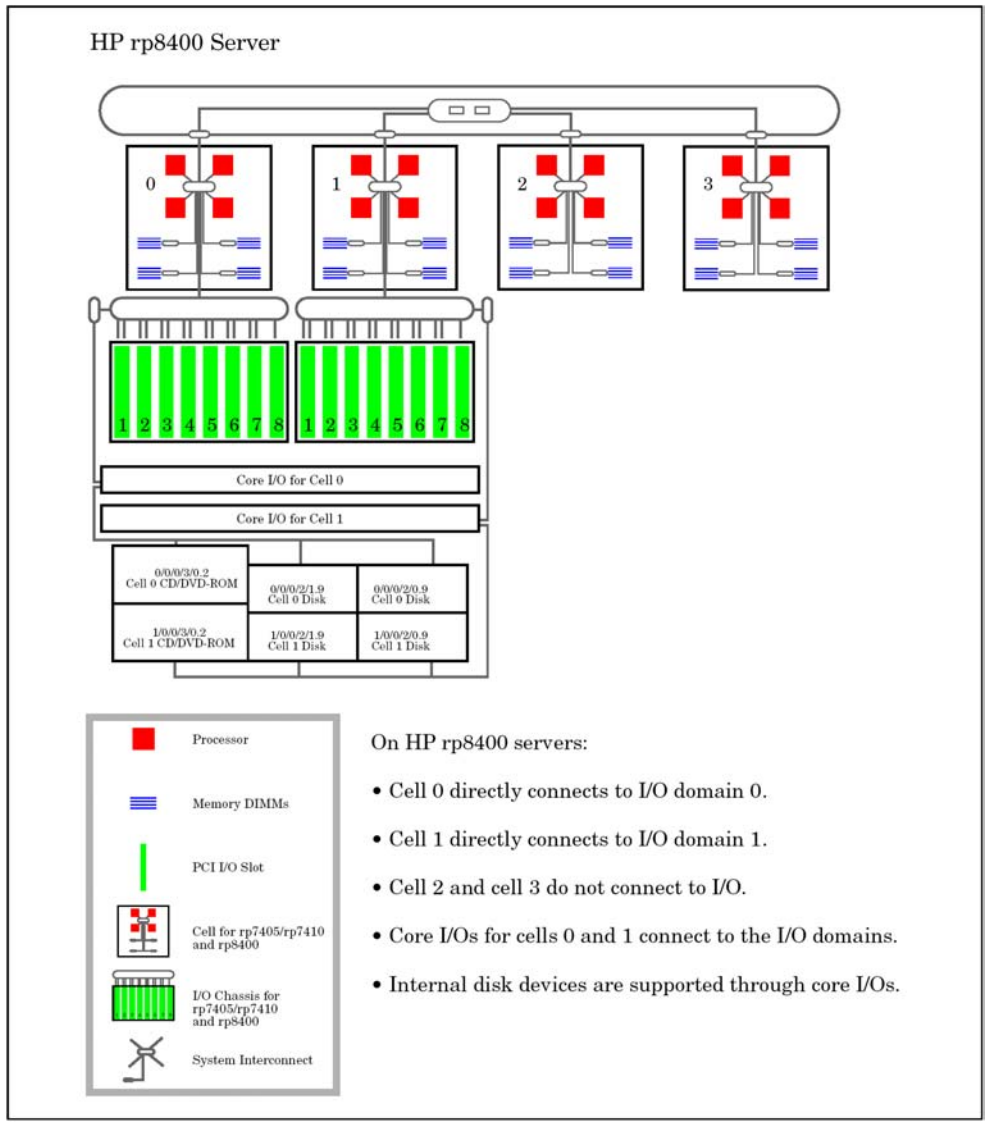

Superdome 32 way nPartitions architecture

HP Superdome 32-Way Server (SD32000)

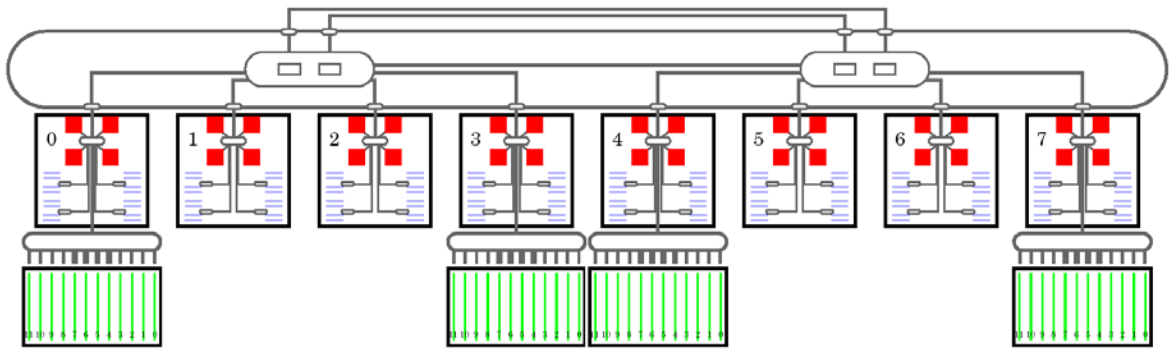

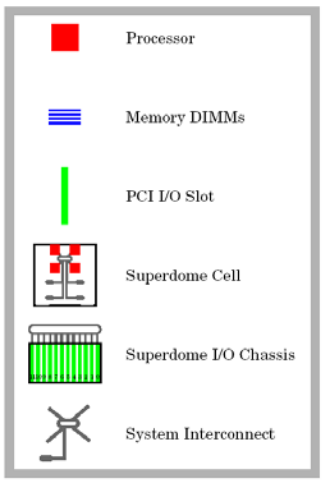

On HP Superdome 32-Way servers:

- Each cell (0–7) can connect to any one of the available I/O chassis.
- Additional I/O chassis can be provided in a connected I/O expansion cabinet.
- PCI card slot 0 in each I/O chassis is for use by a Superdome core I/O card.

HP-UX 11i Virtual Partitions (vPars):

Multiple O/S instances per nPartition with software & resource isolation, and flexibility

- **Increased system utilization**
  - Up to 80-90%, since additional partitions can use unused parts of system
- **Increased isolation**
  - Of OS, applications, resources
  - Individual reconfiguration & reboot
- **Greater flexibility**
  - Multiple independent OSs
  - 1 CPU granularity per vPar
  - Dynamic movement of CPU power between vPars
  - Resources not tied to physical configurations (like hard partitions)
  - vPars on low to high-end servers

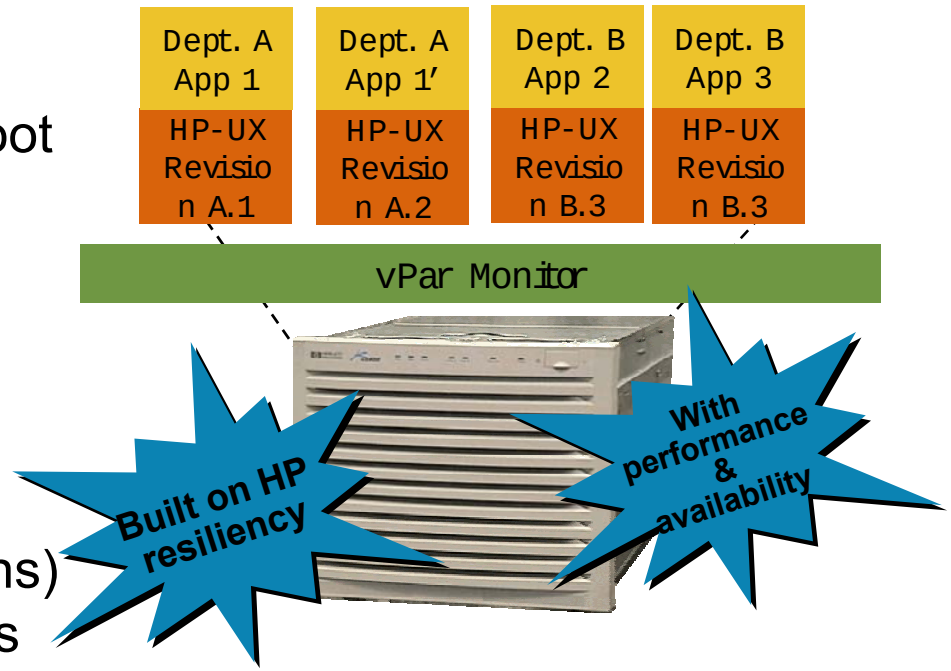

vPars strengths

- Performance: < 5% overhead
- Easy management:
  - vparmgr GUI
  - Automatic, SLO-based workload management ACROSS vPars (WLM cross-vPar integration)
- Easy deployment:
  - Ignite/UX is vPar-aware
- Flexibility:
  - Dynamic CPU migration
  - 1 CPU granularity
  - Integrates with nPartitions and iCOD
  - Allows app-specific O/S tuning
- Resource isolation: resources are assigned to a vPar
- Platform support (as of version A.02.02):
  - rp5405, rp5470/L3000, rp7400/N4000, rp7410, rp8400, Superdome

vPars logical overview

- Multiple applications or multiple instances or versions of the same application
- No name space or resource conflicts

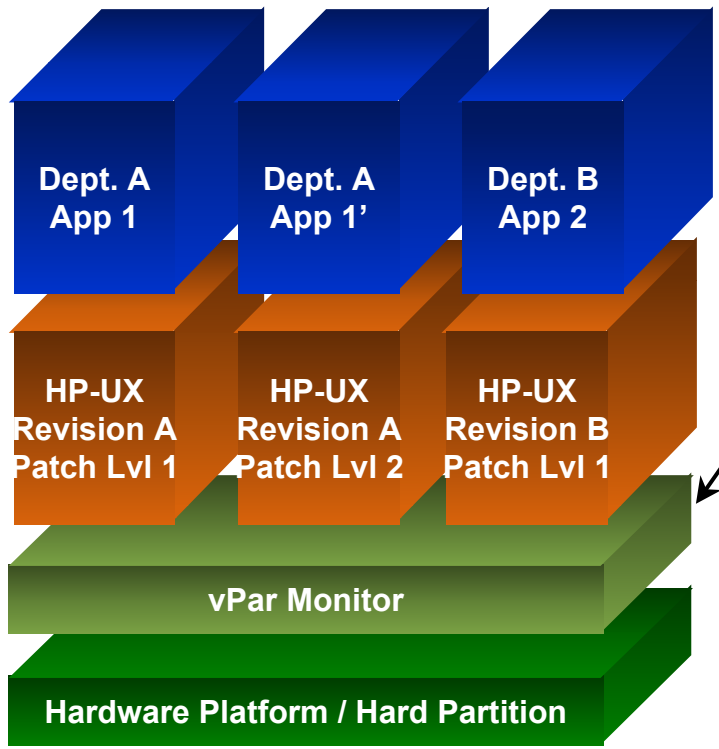

- Creates illusion of many separate hardware platforms
- Manages shared physical resources
- Monitors health of operating system instances

- Each operating system instance tailored specifically for the application(s) it hosts
- Operating systems instances are given a user-defined portion of the physical resources
- No name space or resource conflicts

vPar monitor: Between HP-UX 11i (O/S) & firmware

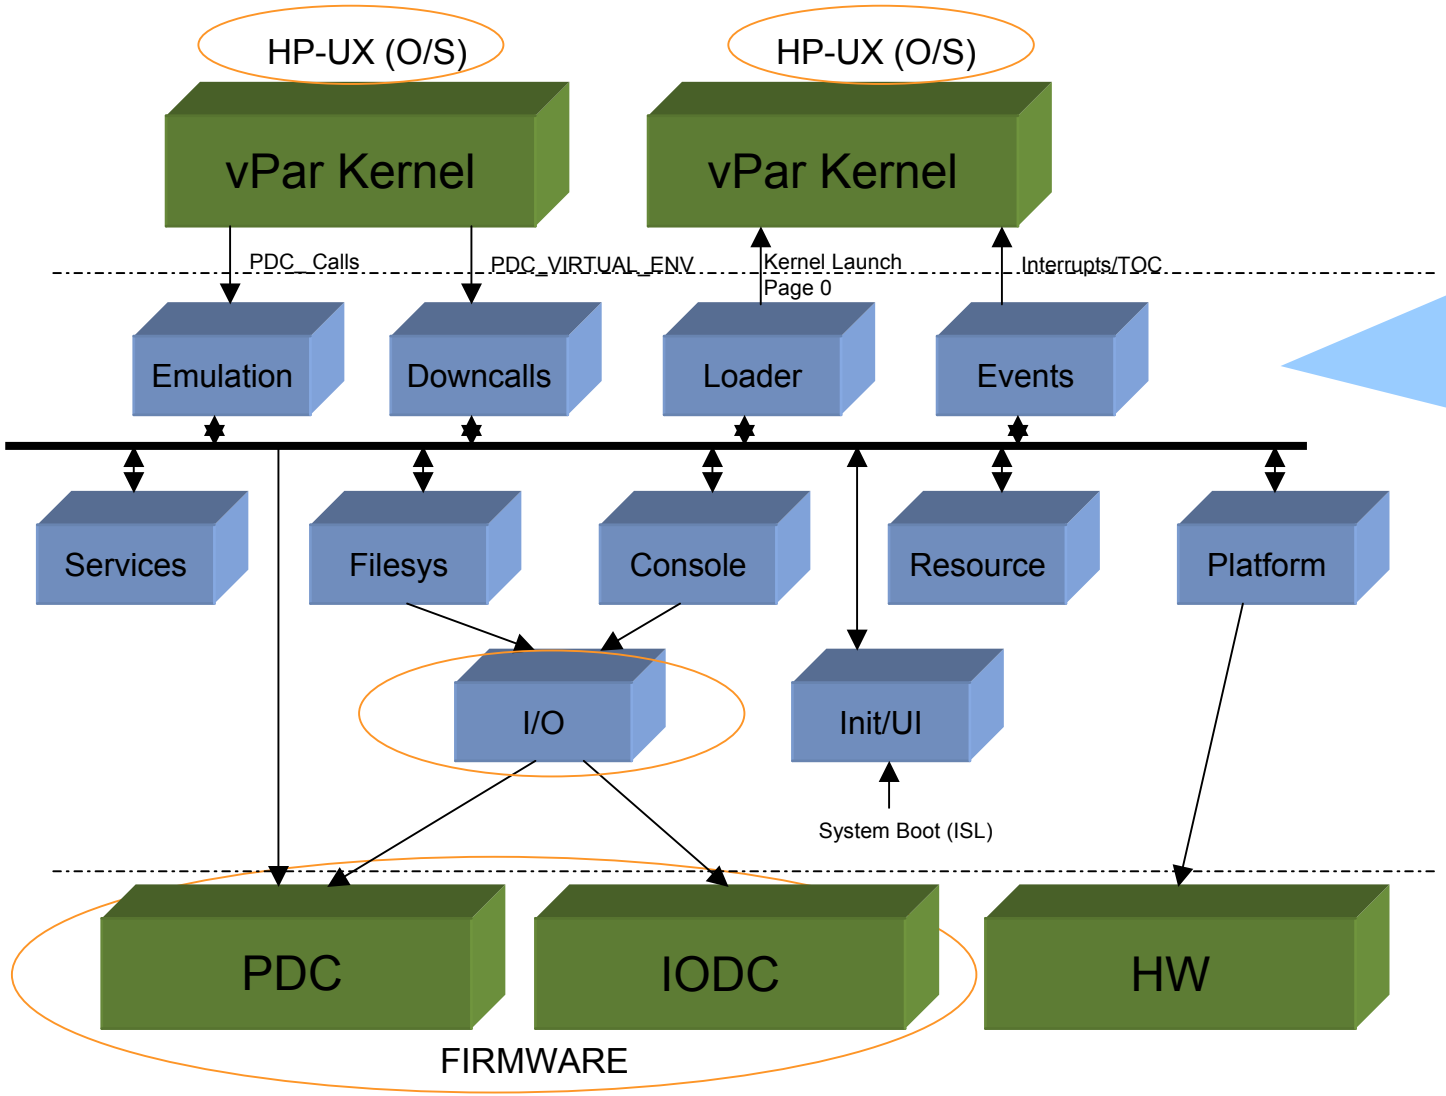

- Minimal performance overhead
- Invoked (mainly) at startup/ shutdown, and during firmware or vPar commands

Memory allocation for vPars

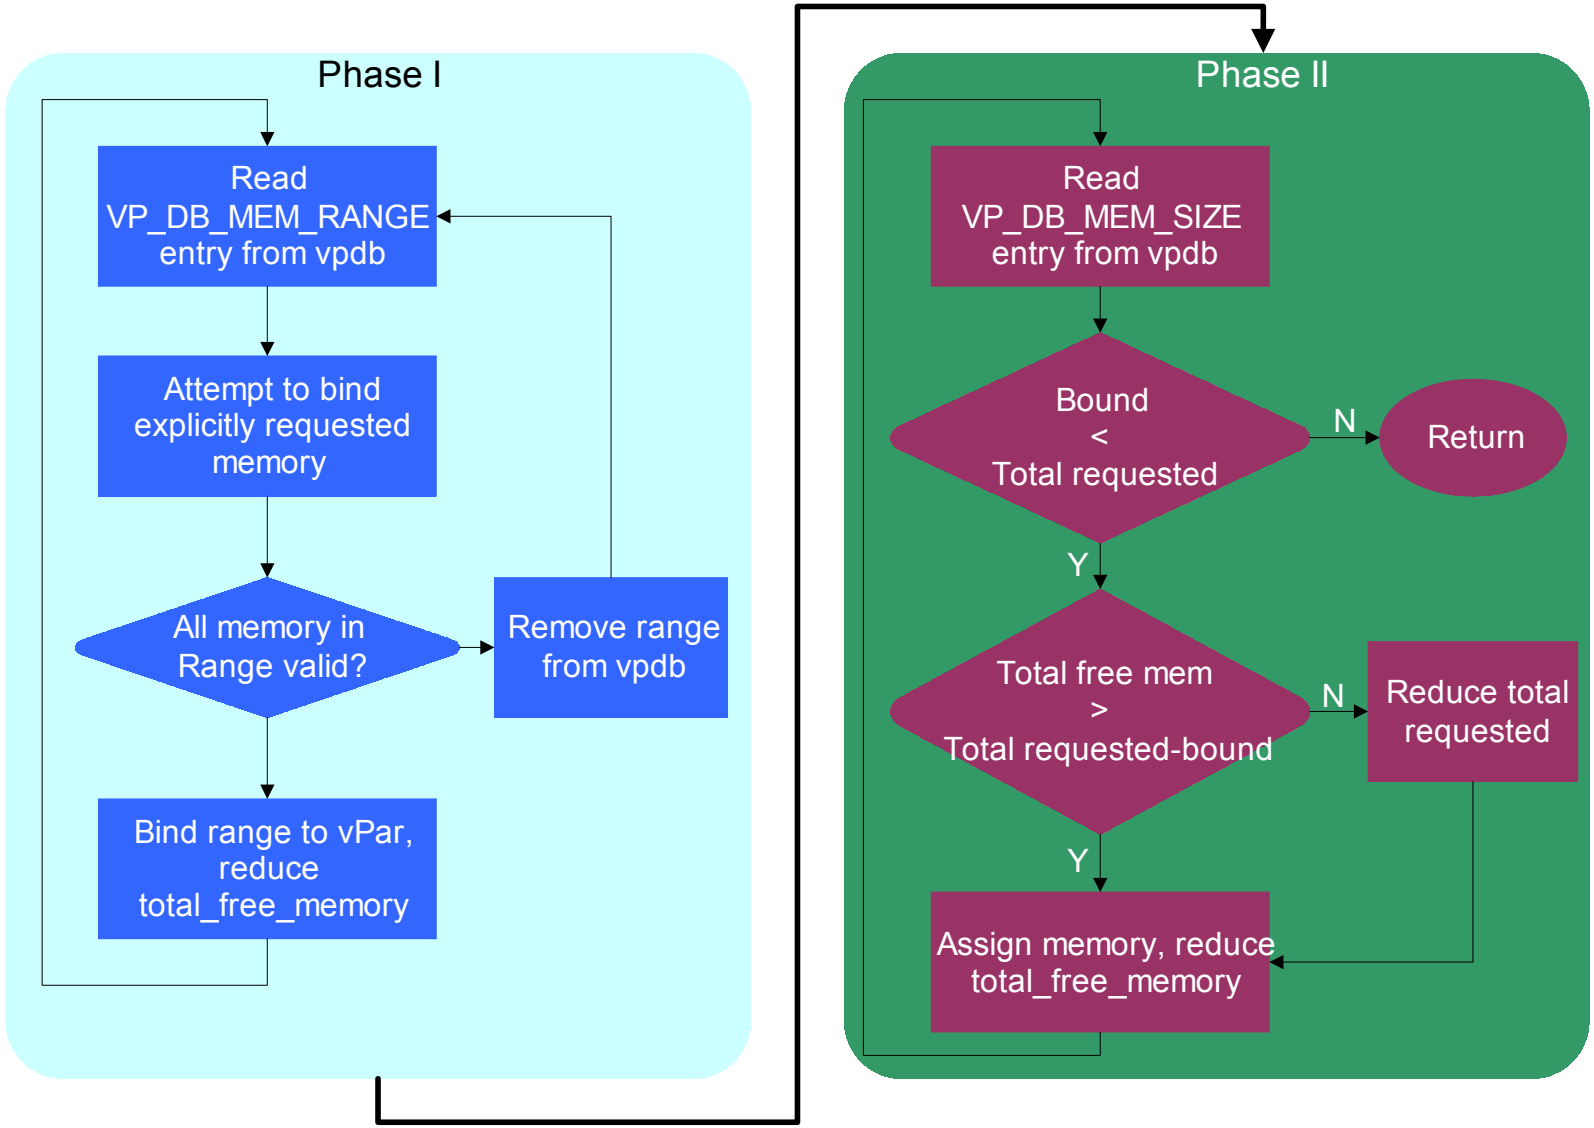

vPar Security, Availability and Performance

Can a vPar steal resources from another vPar?

No, the PDC emulator prevents that

Can a vPar crash another vPar(s)?

No, only a HW fault will affect multiple vPars

Does vPars affect each others performance?

No, since they do not share resources

Will a move of CPU affect the Memory allocation?

No, since HP has a separate memory controller

HP Process Resource Manager (PRM)

Predictable service level management

- Resource partitions within a single OS image

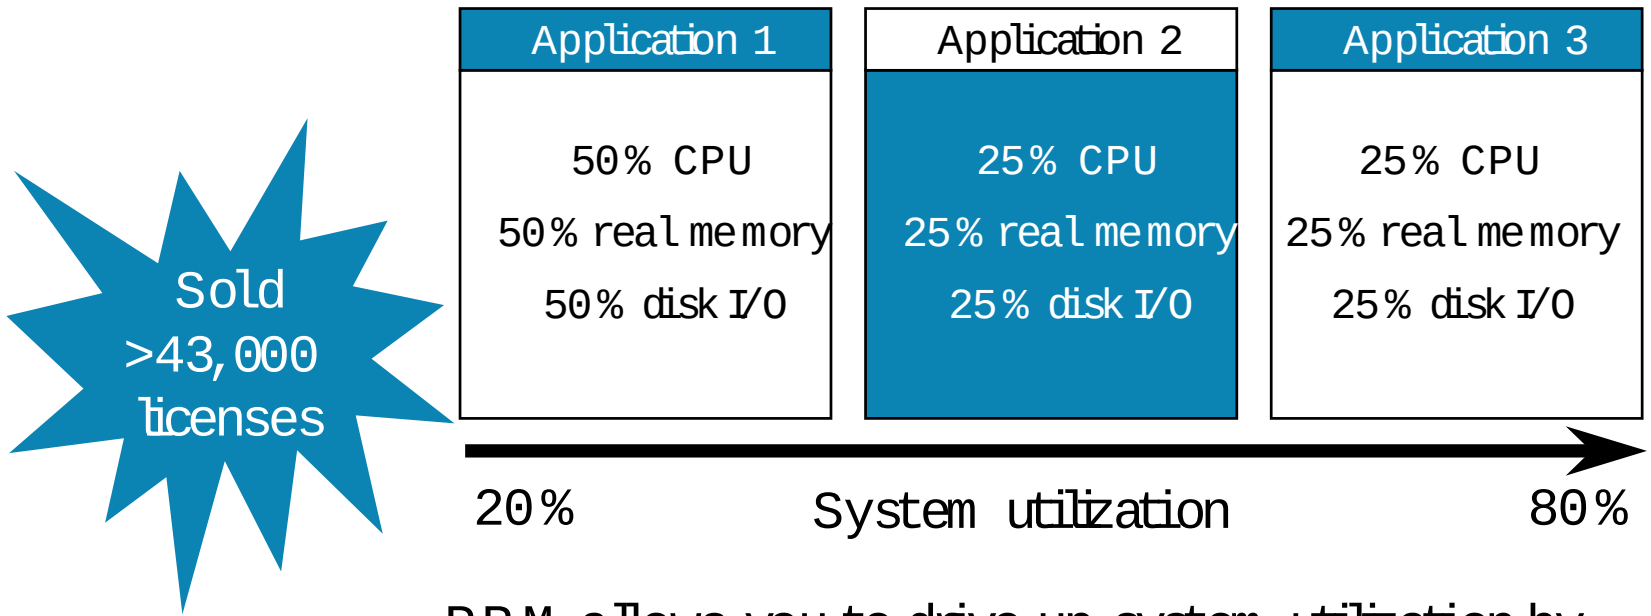

PRM allows you to drive up system utilization by running more applications per server: the result is a better ROI

HP Process Resource Manager features

- Hierarchies (PRM 1.08)
  - Shares support in addition to percentage based allocation (PRM 1.08)
  - In-kernel memory (PRM 1.08)
  - Single-point administration, Java-based GUI (PRM 1.08)
  - Increased resource allocation flexibility by supporting processor sets (PRM 2.0)
  - Offers optimum consolidation and dynamic reallocations for Oracle environments - now also with Oracle Database Resource Manager (PRM 2.0)
  - Supports HP-UX 11i and HP-UX 11i Version 1.6 (PRM 2.0)

HP pSets and PRM

- pSets provide
  - Grouping of CPU's to control processes resource usage
  - CPU resource isolation for applications and users
  - can run within a nPar or vPar
  - Is free with HP-UX 11i (as additional patch)
  - dynamic creation, deletion, and reconfiguration of pSets
- Integration with PRM provides
  - configuration is maintained across reboots
  - can further partition a pSet using the fair share scheduler (FSS) within a pSet
  - Provides memory management within a pSet
  - Automation of processes and user assignment (to pSets)
  - Java based GUI for configuration of psets

Integrated hierarchical mgmt. of Oracle on HP-UX

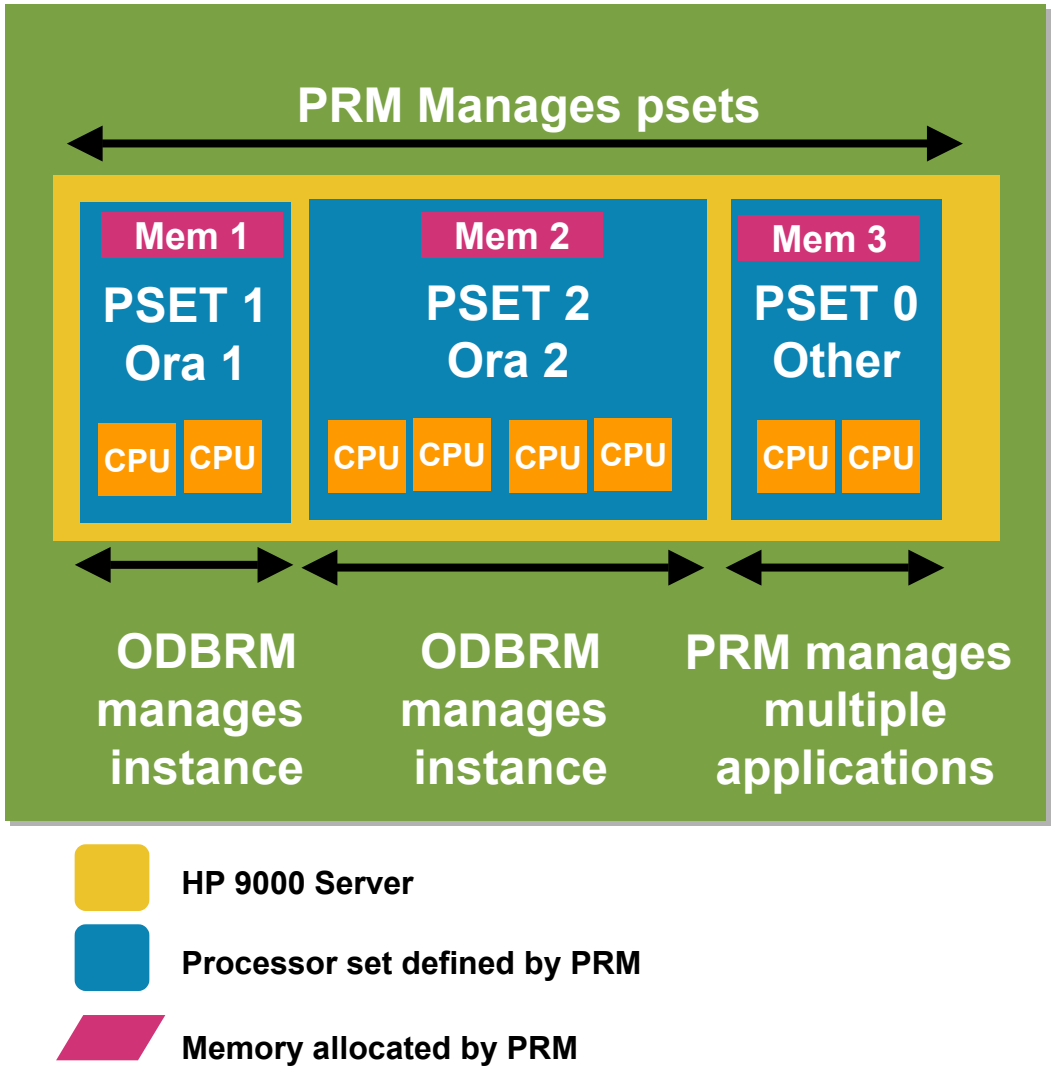

PRM co-functions with Oracle 9i Database Resource Manager (ODBRM) to manage system resources

Result: HP-UX is the ideal Oracle consolidation platform

HP-UX Workload Manager (WLM)

The goal-based policy engine of the HP Virtual Server Environment

Dynamic resource optimization

Automated and intelligent management

Examples of Service Level Objectives (SLOs)

Priority 1	Priority 2	Priority 3
Response time SLO	**Response time SLO**	**Job duration SLO**
Transactions will complete in less than 2 seconds	**Transaction will complete in less than 3 seconds**	**Batch job will finish in less than 1 hour**
Application a	**Application b**	**Application c**

Automatic reconfiguration of CPU resources to satisfy SLOs in priority order

Combining the partition continuum

More flexibility, isolation & granularity

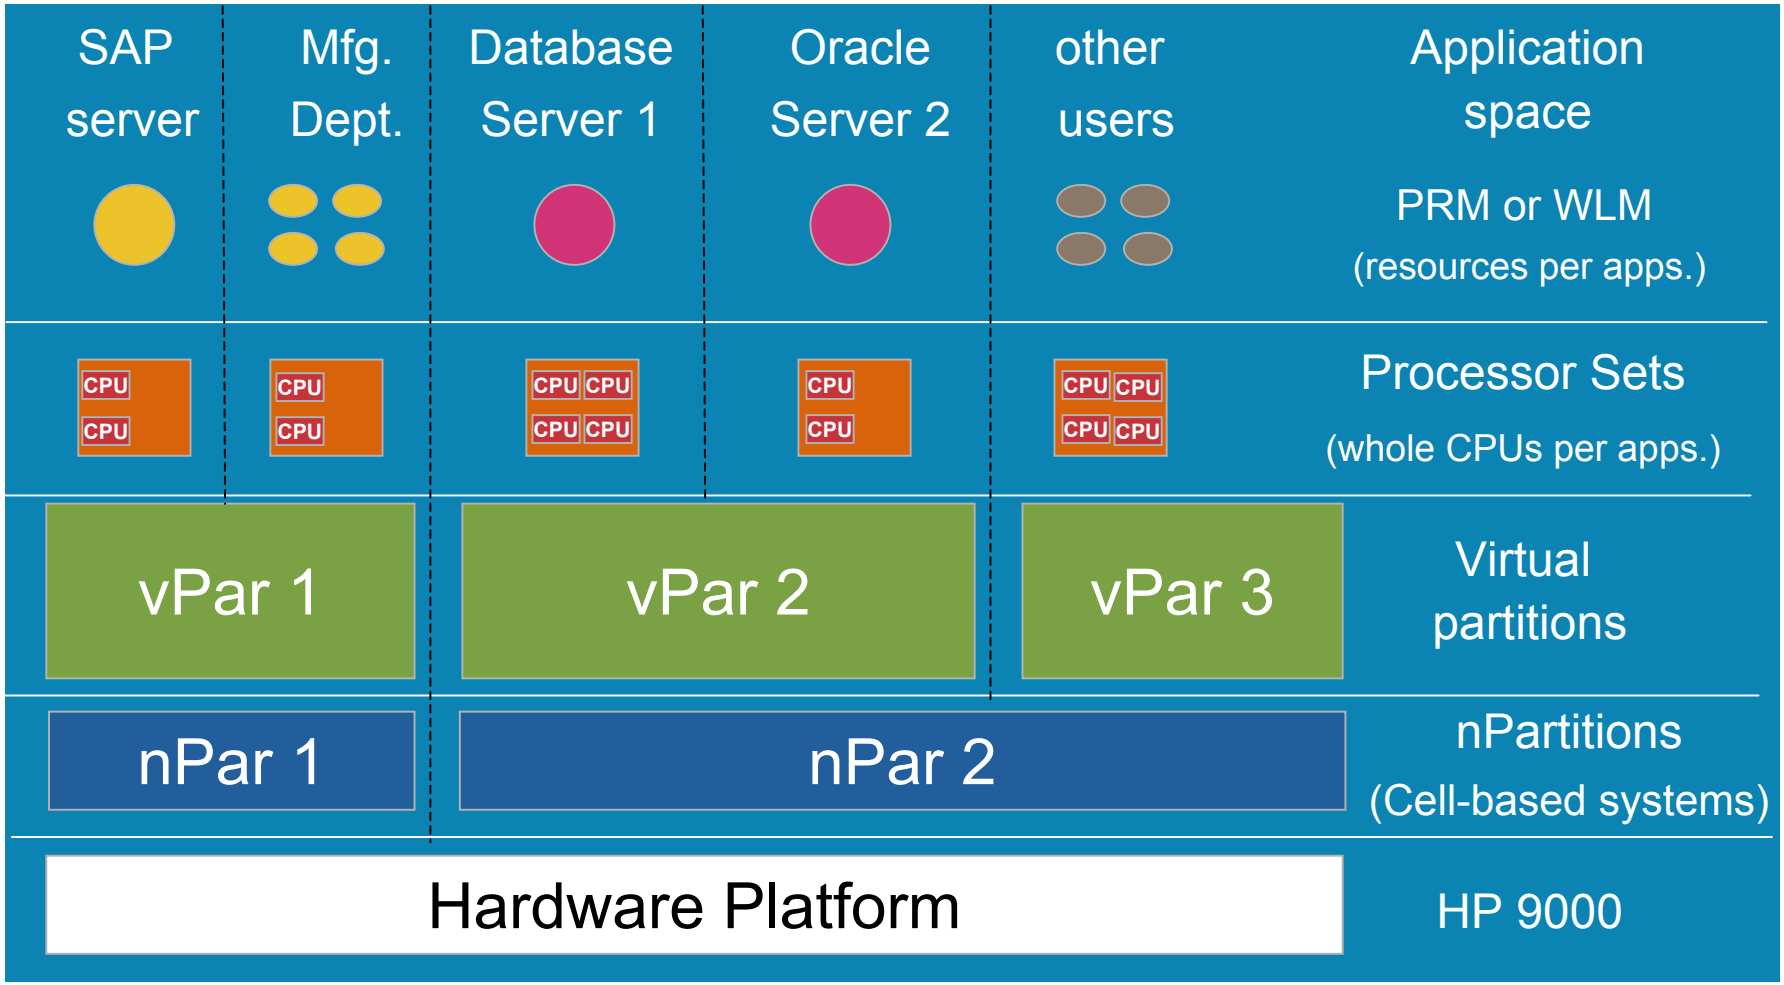

HP-UX server virtualization: WLM and Serviceguard

HP WORLD 2003 Solutions and Technology Conference & Expo

Automatic resource adjustment upon failover

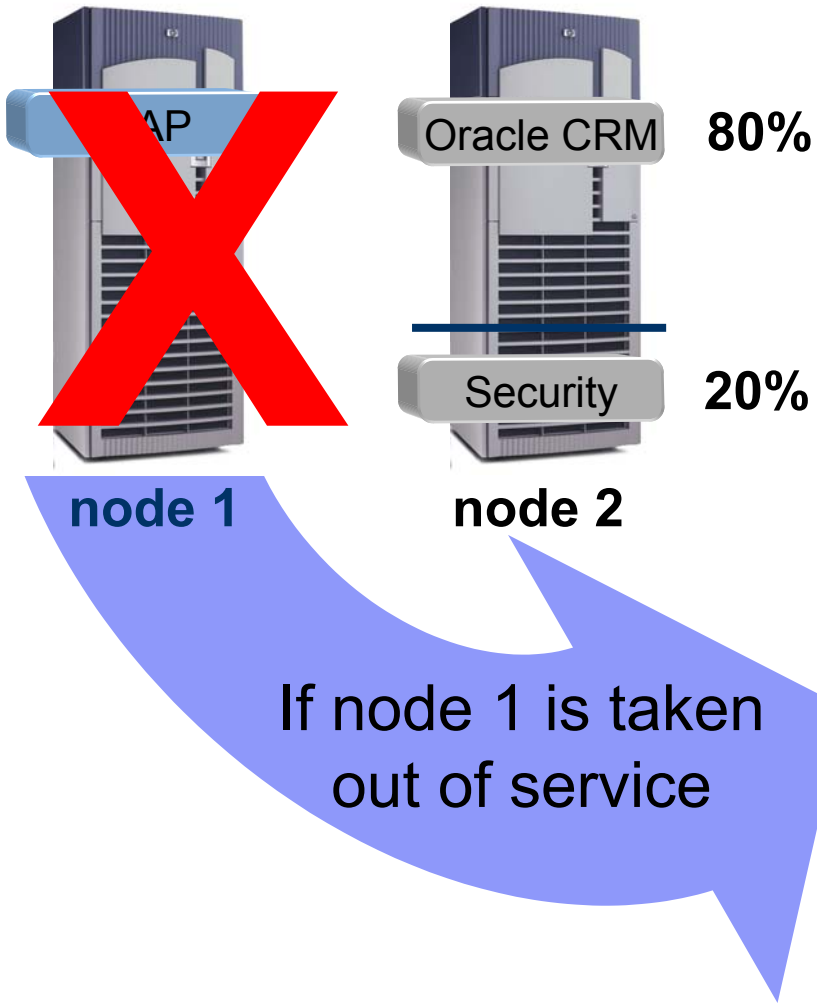

- Customer sets business priorities for each application
- HP Virtual Server environment automatically ensures the fulfillment of business priorities despite server downtime
- Dynamic reallocation of system resources
- Load balancing for normal and post-failure operation

HP-UX server virtualization: vPars, WLM & Serviceguard

Automatic vPars adjustment upon failover

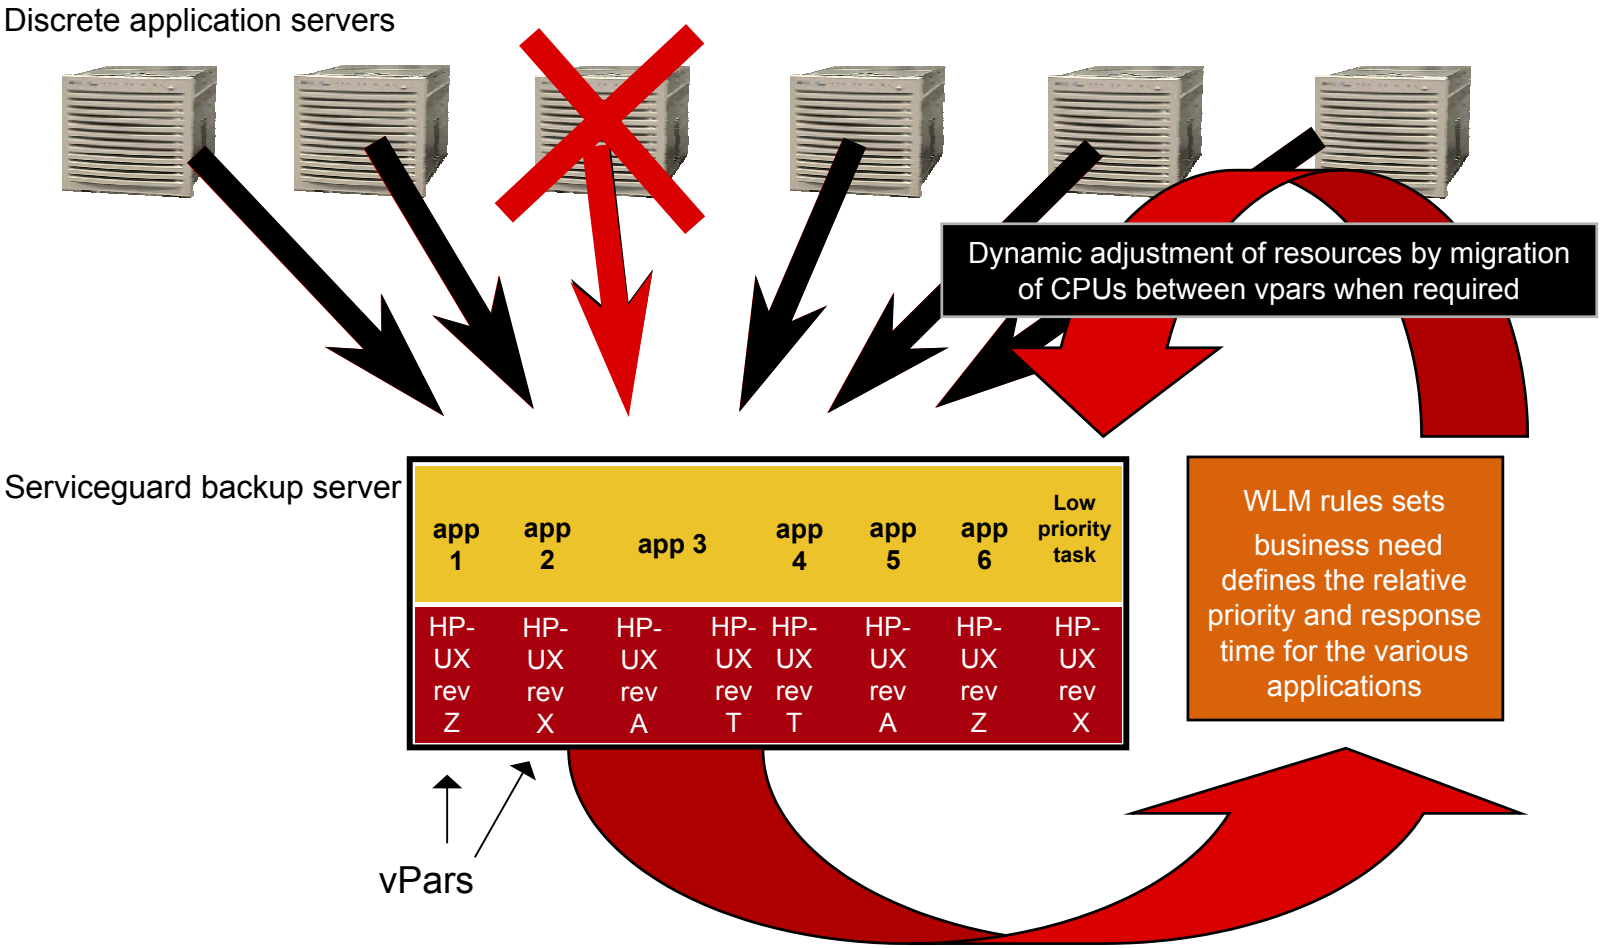

HP-UX Server virtualization:

ServiceGuard, WLM & iCOD
Automatic resource adjustment & iCOD enablement, upon application migration

Scenario: Customer needs to dedicate node 1 to end of month financials (App A).

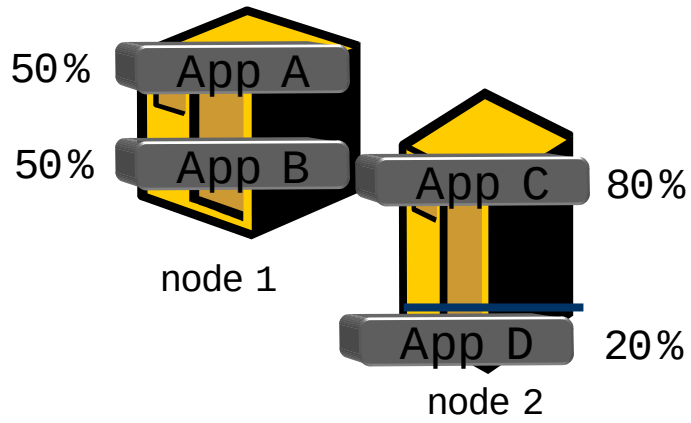

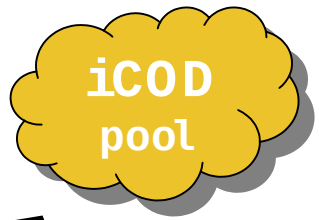

1) **Serviceguard** migrates App B to node 2

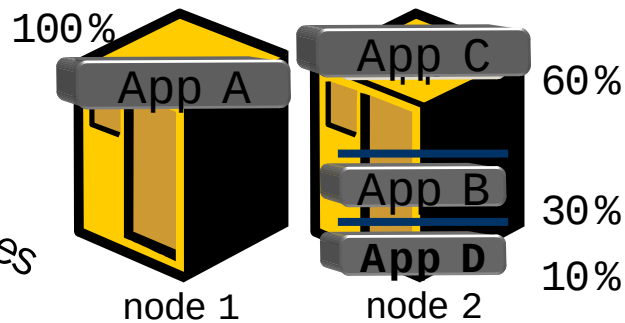

2) **WLM automatically:**

a) Is aware of new application (App B)

b) Adjusts application resources based on pre-defined SLOs

c) Notifies or enables additional iCOD as needed

Server consolidation Ex: Dev & test nPars with vPars

"Right" level of application isolation with uptime

- nPartitions provide electrical hardware isolation - separate systems, I/O, boot, CPU, memory, etc.
- vPars are software isolated (OS, middleware, apps) for different customers, different OS tuning - each tailored by number of CPUs, amount of RAM, amount of I/O per customer and application mix

•Electrical H/W isolation

•S/W isolation
•Dynamic reconfiguration
•Single CPU granularity
•Low to high servers
•Resources not tied to physical configuration

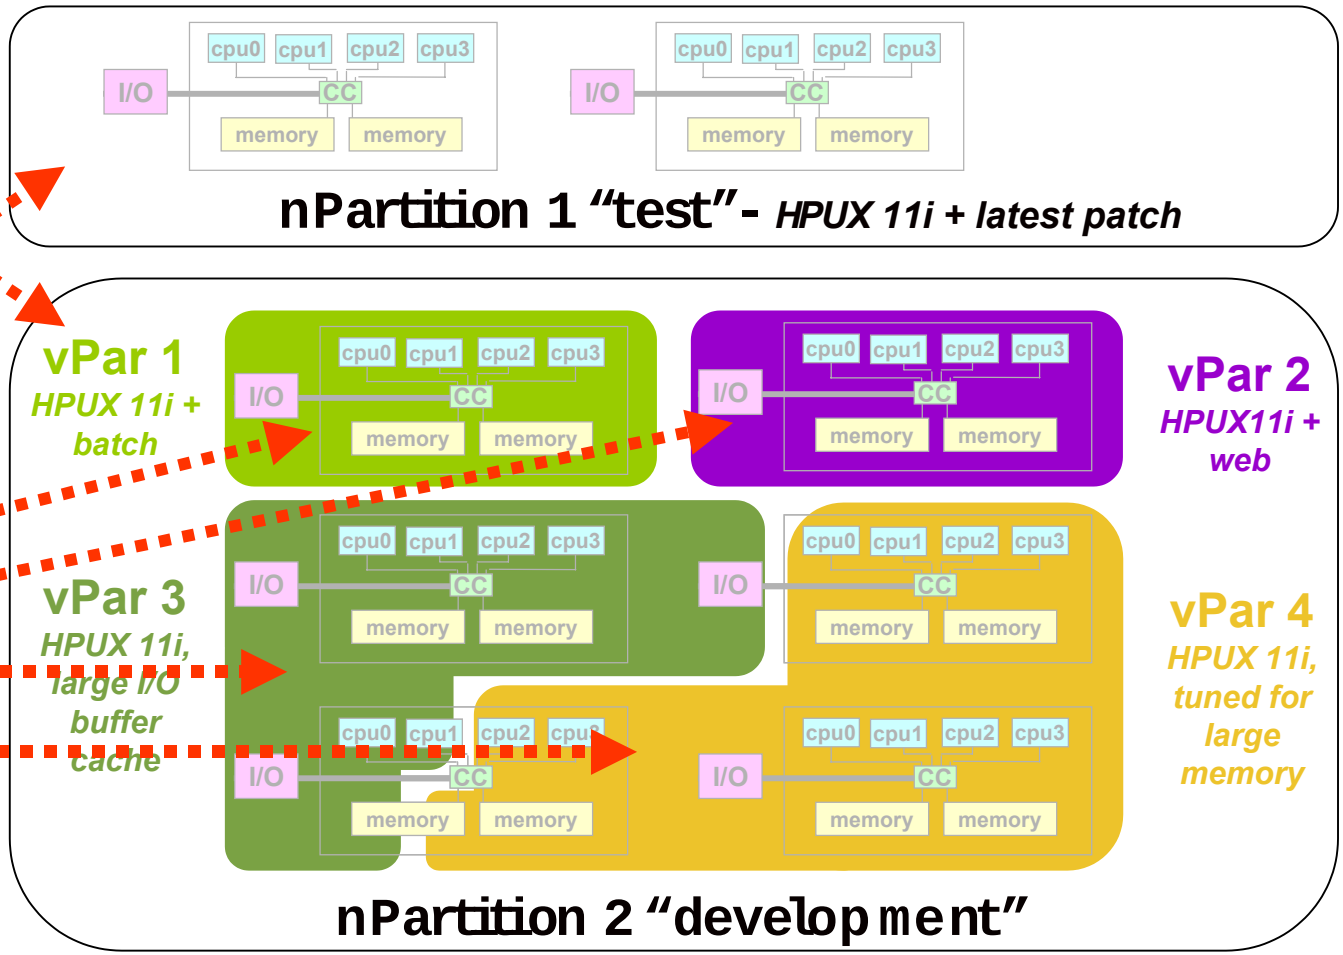

Managing the HP partitioning continuum

The integrated power of HP ServiceControl manager and HP OpenView

ServiceControl manager

- Single-point multi-system management
- Management for rapid deployment and consistency

Partition manager

- Sets up partitions
- Maps partition architecture

HP OpenView GlancePlus pak

- Monitors performance of each partition

HP OpenView VP Operations

- Monitors events on each partition

Managing the HP-UX partitioning continuum

Power of HP ServiceControl Manager and HP OpenView

nPartition 1
for test environment

nPartition 2
for production environment

vPar 1,2,3

vPar 1,2,3

Servicecontrol Manager
- Central, consistent management interface
- Grouping functions
- Role-based security

parmgr & vparmgr GUIs
(or command line) create, configure & map partitions

Ignite/UX
for rapid deployment

HP-UX workload manager
automatic resource management by pre-defined SLOs

Glance Plus Pak
performance monitoring of partitions

OpenView operations
monitors events on each partition

HP-UX 11i reduces management complexity

HP WORLD 2003
Solutions and Technology Conference & Expo

Automated and intelligent management

Servicecontrol Manager

Integrates HP-UX and 3rd party tools

Executes multi-system commands simultaneously

Reduces operator error via role-based security

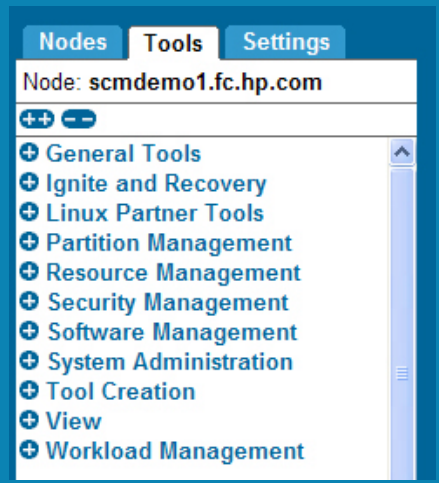

Rapid deployment (Ignite/UX)

Deploys system images and packaged software

System inventory manager

Asset management

Snapshot for troubleshooting

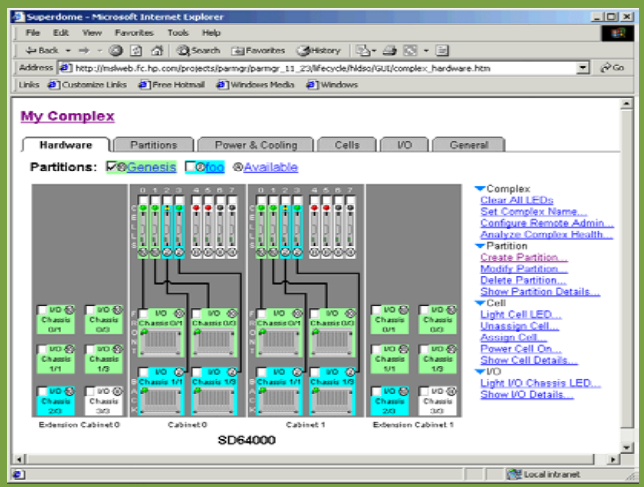

Partition Manager

Configures partitions

Real-life, big picture view for hard-partitioned servers

(1st on 11i v2; later on 11i v1 & w/ vPars)

Workload Manager

Adjusts resources automatically based on predefined SLOs

And many more…

The near future

HP first with multi-system intelligent policy engine (gWLM)

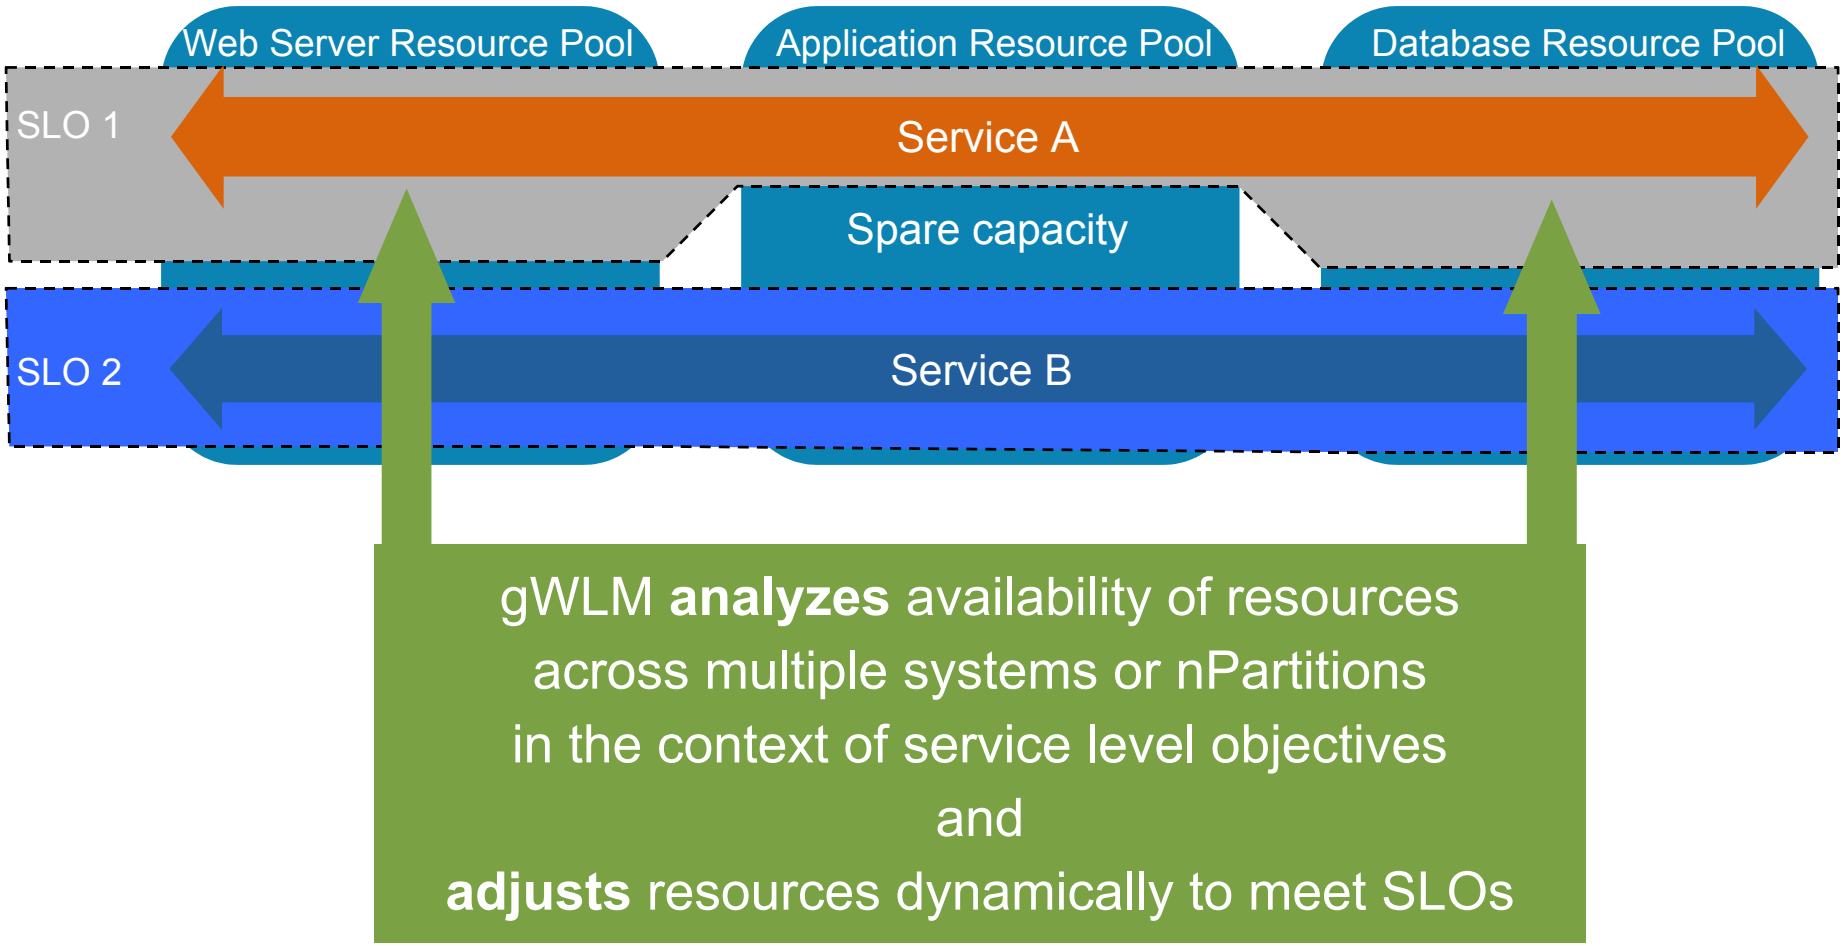

Server centric view of UDC and gWLM

UDC today

- Activates servers to create a new service
- Monitors server performance and flexes based on overall server utilization

gWLM future

- Monitors service performance and reassigns resources within the servers to meet SLOs
- Requests that the UDC flexes servers to meet SLOs

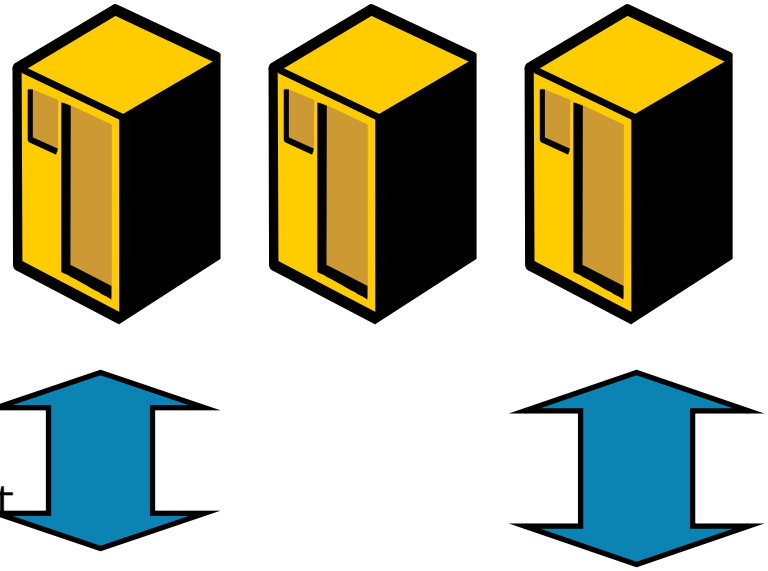

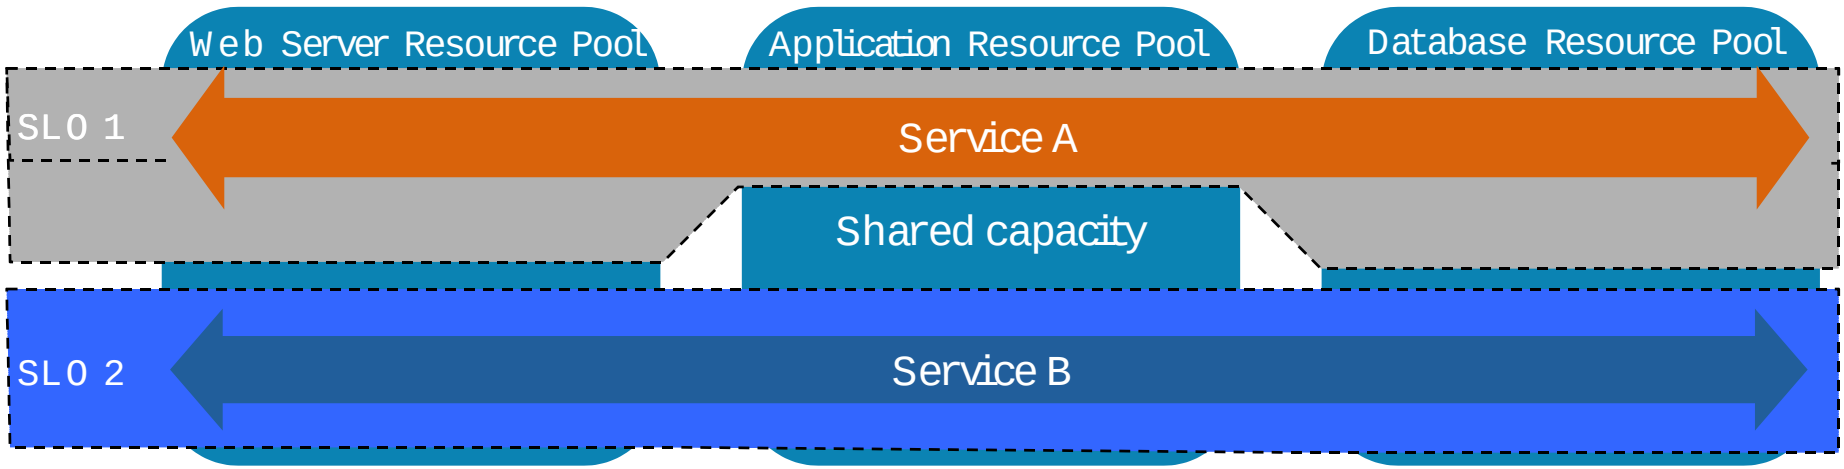

Case Study #1 HPShopping.com

Problems with previous environment

- Customer: HPshopping
- Industry: e-commerce/web retail
- Previous Environment
  - Production and Disaster Recovery in separate data centers. DR under-utilized and expensive to maintain.
  - Little flexibility with hardware resources e.g. new test environment needed.
  - Lots of mid range servers – expensive to purchase and maintain.
  - Applications intermingled e.g. J2EE applications with BroadVision presentation layer. Implications for differing application requirements, performance and O/S level changes.
  - Different tiers having imbalance in utilization e.g. web server having resources that are underutilized, database over utilized.
  - Environment could not scale to meet the projected growth e.g external storage would need to be replaced.

Case Study #1 HPShopping.com
Benefits of a partitioned solution

- Server consolidation from 35 mid-range servers (N/L class) to 3 SuperDome's (and A-class web servers)
  - 2 – 32 way SuperDomes in Production
    - Production SD with 12 nPars and 8 vpars
  - 1-32 way SuperDome in Development / Test / Staging
    - Development SD with 14 OS images
    - 5 nPars and 12 vPars
    - Flexibility, Lower maintenance and support costs
    - Performance and Scalability
    - Cost savings (1 SuperDome in savings)
- Part of vPars Early Adopter Program (EAP). This meant early access to the software as well as consulting and verification of design. The hpshopping implementation was the 1st implementation internally and one of the 1st on SuperDome.

Case Study #2
Pitney Bowes
Problems with previous environment

- Customer: Pitney Bowes
- Industry: Business Communication
- Previous Environment
  - Older leased HP systems (R-class, K-class, L-class)
  - New Broadvision Development Project
- Requirements
  - Isolate different lines of business environments
  - Create development/test/QA environments
- Reduce response time to Lines of Business (LOBs) for creating and deploying development/QA/test environments from the traditional 6 weeks (for power, network, etc.)
- Be able to re-deploy and re-provision resources as needs changed

Case Study #2
Pitney Bowes
Benefits of a partitioned solution

- Mixed production, development, QA and testing – Used vPars to re-deploy unused capacity from production environments to support other applications.
- Very happy with the capabilities of the SD to speed deployment, re-deploy capacity for better overall utilization, resource sharing and the isolation capabilities of nPars.
- 3 - 32 way Superdome Complexes
  - 21 nPars
  - 23 vPars in 10 nPars
- Running(Development/Test/Training/QA/Production/Serviceguard Clusters)
  - Broadvision, Oracle, Siebel, SAP BW(Business Warehouse), Windchill, Customer Developed Applications
- With the SD deployed new environments can be created within days to support the business instead of weeks

Case Study #3
Financial Services example
Problems with previous environment

- Customer: anonymity requested
- Industry: Financial Services
- Environment
  - New ServiceDesk Implementation Project
  - Multi-tiered solution consisting of Application Server, Web Server, and Database Server
  - Flexible Resource Allocation / Re-allocation based on peak workloads
  - Local ServiceGuard Failover with remote failover capability provided via DB replication.
- 3 production rp8400s, fully populated, with 1nPar and 4vPars each
- 1 test rp8400 with 2 nPars and 3 vPars
- Part of vPars Early Adopter Program (EAP). This meant early access to the software as well as consulting and verification of design. This was the first implementation of vPars on the rp8400

Case Study #3
Financial Services example
Benefits of a partitioned solution

- Rapid system resource allocation or reallocation has cut deployment time from weeks to hours.
- Dynamic resource allocation to meet peak system processing needs
- Was able to maximize performance, scalability and flexibility while minimizing maintenance and support costs and system footprint.
- Ease of support of “monoculture systems” (all OS’s versions and patch levels are identical across all vPars)
- Improved overall system resource utilization by having resources provisioned where and when needed (Just in Time Resource Delivery)

Customer Quote *“ our confidence in HP, their vPar product, and their ability to support the solution gave a sense of security to our deployment of this cutting edge technology”*

White board example

Summary

- HP-UX has the only goal-based workload management (intelligent policy engine) in the UNIX industry
- HP-UX has the broadest virtualization capabilities in the UNIX industry: (Hard and soft partitions, resource partitions with processor sets)
- HP continues its leadership in multi-OS system management with Servicecontrol manager, (and its planned integration with Insight Manager)
- Global Workload Manager is the future multi-system workload management solution

Questions?

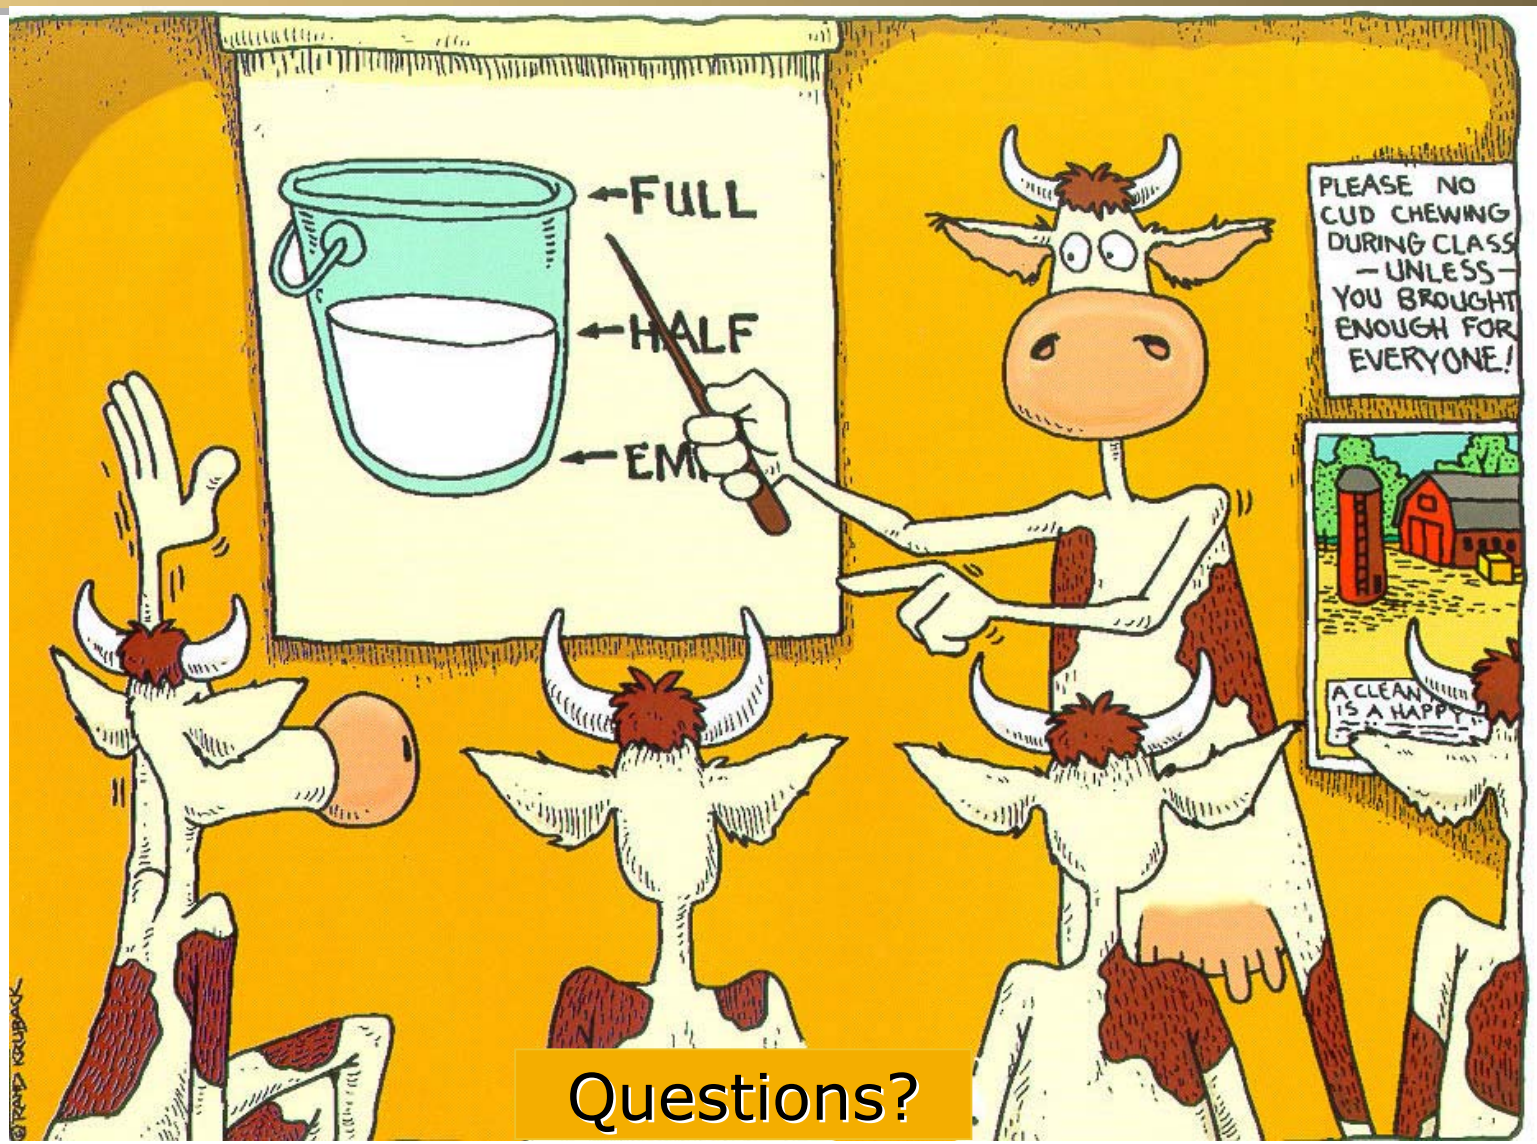
FULL
HALF
EMP
PLEASE NO CUD CHEWING DURING CLASS —UNLESS— YOU BROUGHT ENOUGH FOR EVERYONE!
A CLEAN IS A HAPPY
Questions?

Interex, Encompass and HP bring you a powerful new HP World.

Backup

HP-UX Workload Manager –
the intelligent policy engine ties business priorities to resources

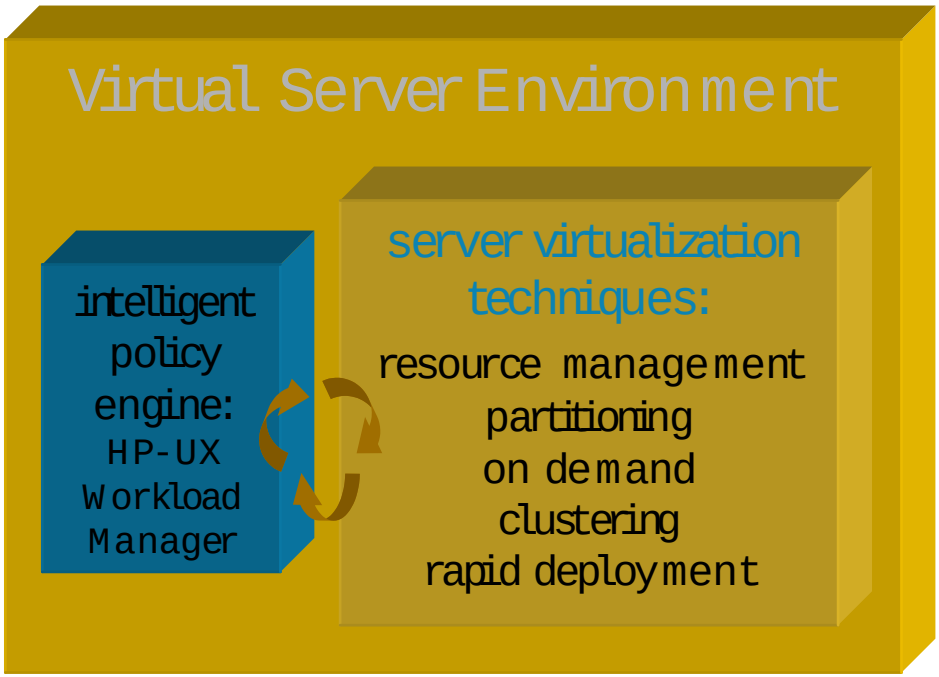

- Goal-based resource management – automatic allocation of CPU resources based on set service level objectives and business priorities
- Predictable response times for mission-critical applications
- Process Resource Manager included
- Support of advisory mode (*)
- Application transparency
- Support of Application Response Measurement (ARM)

Out-of-the box toolkits

- Oracle database toolkit
- Apache toolkit
- SAS software toolkit
- BEA WebLogic toolkit (*)

Unique integration with virtualized server resources to optimize system utilization!

(*) new functionality with WLM 2.1 – June 2003

HP Partitioning Continuum Products across HP OS's

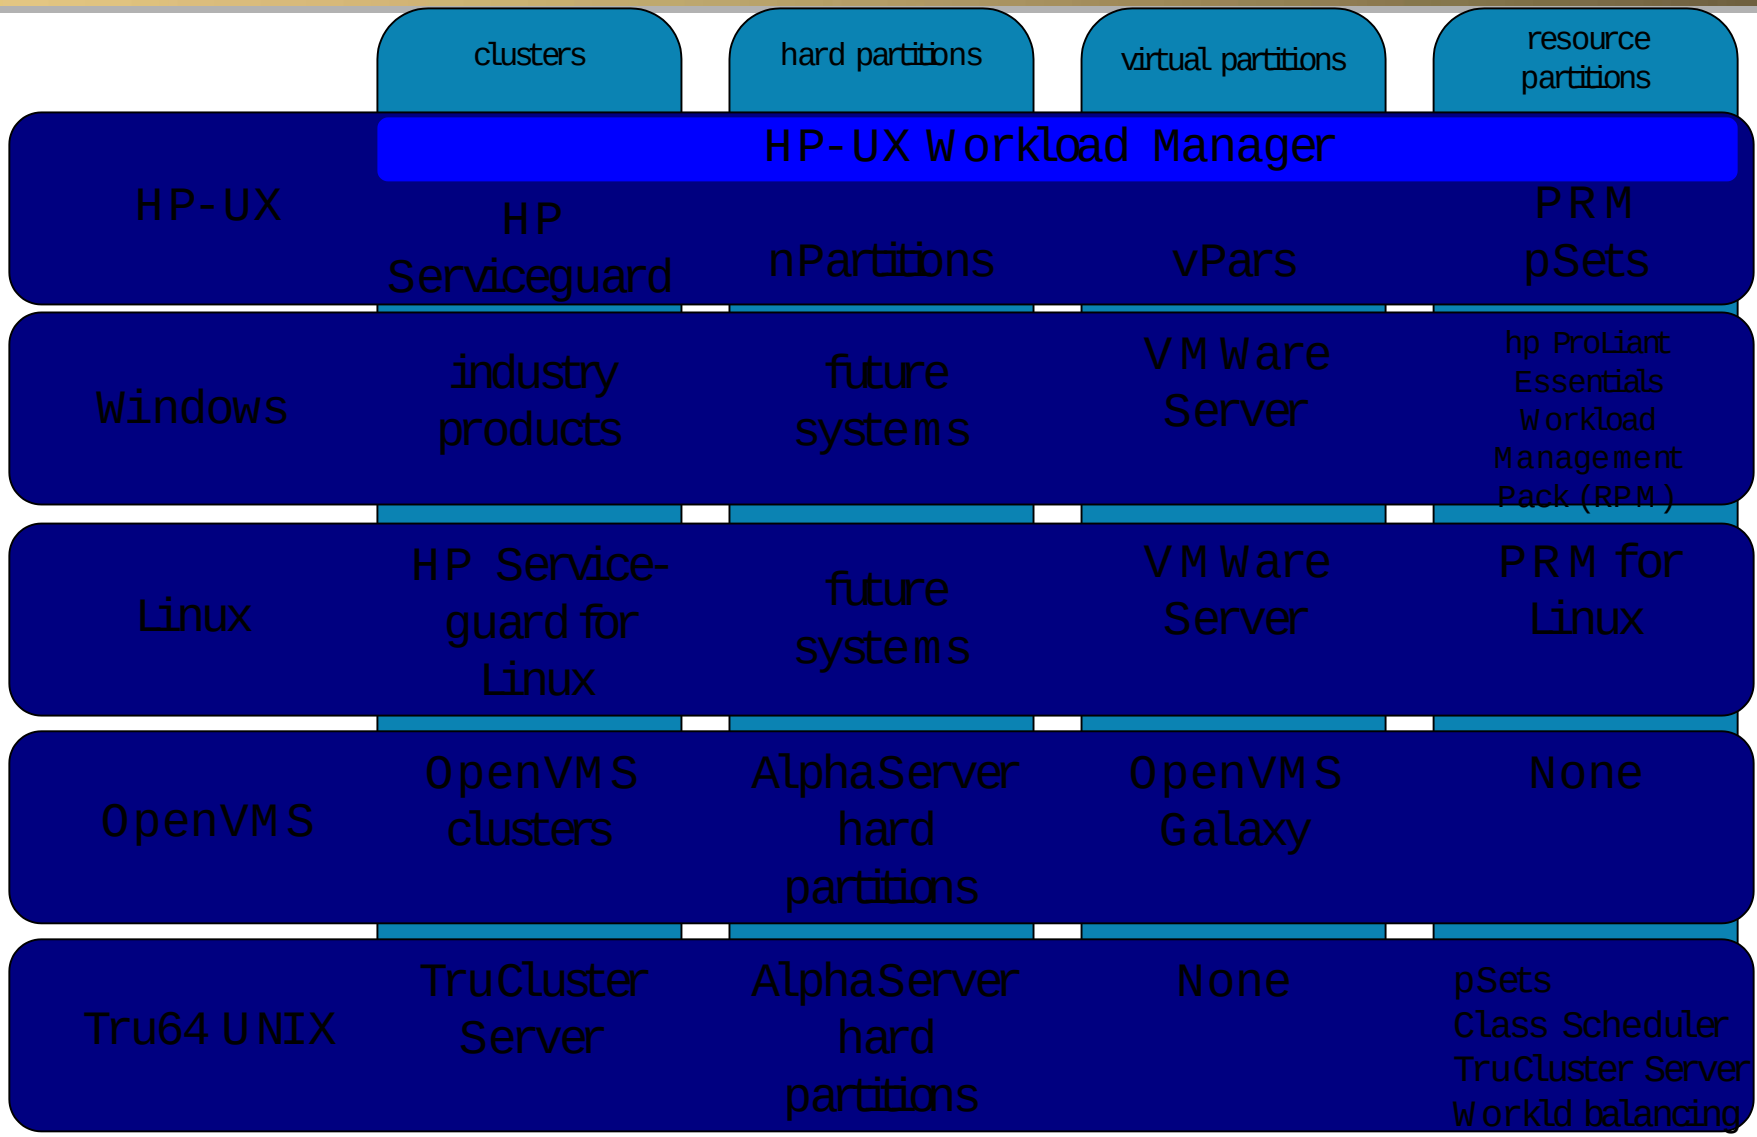

HP: a better approach

- Offers best RoIT
- Leader in high availability, manageability and virtualization
- Delivers best-in-class solutions for adaptive infrastructure
- Leverages strong partner relationships
- Provides customers with technologies that change ROI (i.e., UDC, Linux, OpenSAN, ZLE…)
- Begins with IT consolidation to establish best business practices