

Improving Utilisation and IT Agility Through Partitioning and Workload Management



Ross Templeton

Operating Environments Marketing
Manager

Asia Pacific Business Critical Systems

Agenda



- Partitioning Overview
- Hard Partitions
- Virtual Partitions
- Resource Partitions
- Workload Management
- Customer Examples
- Demonstration

Definition of partitioning



Partitions are physical or logical mechanisms for **isolating operational environments** within single or multiple servers to offer the **flexibility of dynamic resizing** while ensuring that applications can enjoy **protection from unrelated events** that could otherwise cause disruption, interruption, or performance degradation.



Why is partitioning important?



Why is partitioning important?

- Pressure to offer service level guarantees at reasonable costs
- Under utilisation of servers
- Address high fluctuation of web and application traffic
- Flexibility with privacy and high availability



HP partitioning customer benefits

- Meet service level agreements with best return-on-investment
- 80-90%+ utilisation of compute power
- Fast and dynamic implementation of changing requirements
- “Right” level of application isolation with uptime

HP Partitioning Continuum products



	clusters	hard partitions	virtual partitions	resource partitions
Windows	industry products	future systems	VMWare server	HP ProLiant Essentials Workload Management Pack (RPM)
Linux	industry products	future systems	VMWare server	PRM for Linux
OpenVMS	OpenVMS clusters	AlphaServer hard partitions	OpenVMS Galaxy	None
Tru64 UNIX	TruCluster server	AlphaServer hard partitions	none	ARMTech psets class scheduler
HP-UX	Hyperplex	nPartitions	vPars	PRM pSets
HP-UX Workload Manager				

HP-UX partitioning continuum technical positioning



Hard Partitions
with multiple nodes

Hard Partitions
within a node

Virtual Partitions
within a hard partition

PRM with psets
resource partitions within
a single OS image

<p>HP cluster</p>	<p>nPartitions</p>	<p>Virtual partitions</p>	<p>PRM (Process Resource Manager)</p>
<ul style="list-style-type: none"> – complete hardware and software isolation – node granularity – multiple OS images 	<ul style="list-style-type: none"> – hardware isolation per cell – complete software isolation – cell granularity – multiple OS images 	<ul style="list-style-type: none"> – complete software isolation – CPU granularity – dynamic CPU migration – multiple OS images 	<ul style="list-style-type: none"> – dynamic resource allocation – share (%) granularity – 1 OS image

HP-UX workload manager (WLM)
–automatic goal-based resource allocation via set SLOs



Isolation
highest degree of separation

Flexibility
highest degree of dynamic capabilities



Hard Partitions



nPartitions



multiple applications
on the same server
with hardware isolation



- Increased system utilisation
 - partitioning server into physical entities: up to 16 nPartitions
- Increased Flexibility: Multi OS
 - multi OS support: HP-UX, Linux (*), Windows (*)
 - multi OS version support
 - multiple patch level support
- Increased Uptime
 - hardware (electrical) and software isolation across nPartitions
 - Serviceguard support (within server or to another HP 9000 server)
- available on Superdome, rp8400, rp7410


Virtual Partitions



HP-UX 11i Virtual Partitions:



- Multiple O/S instances per nPartition with S/W & resource isolation, and flexibility
 - increased system utilisation
 - 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



Resource Partitions



Resource Partitions



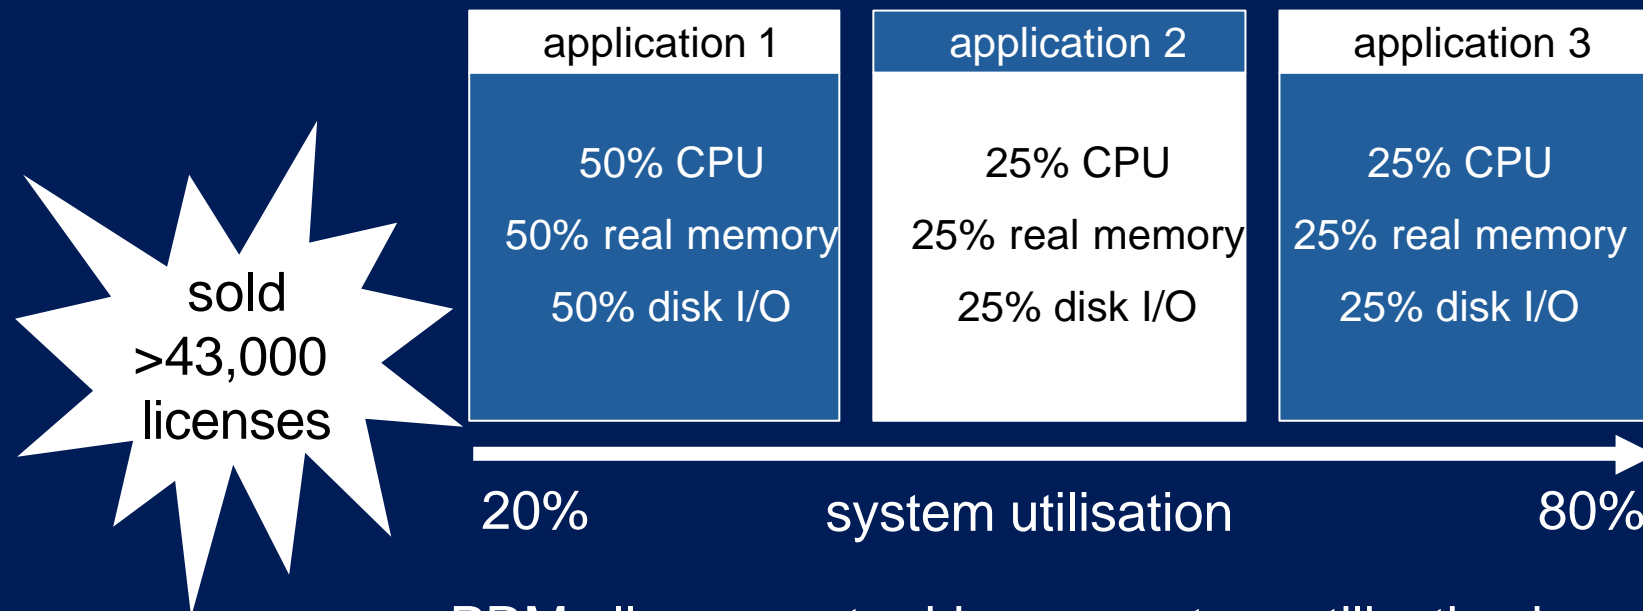
- The Problem: Competition for resources within an OS instance

- The Solution: Resource Partitioning with PRM

HP process resource manager (PRM): predictable service level management

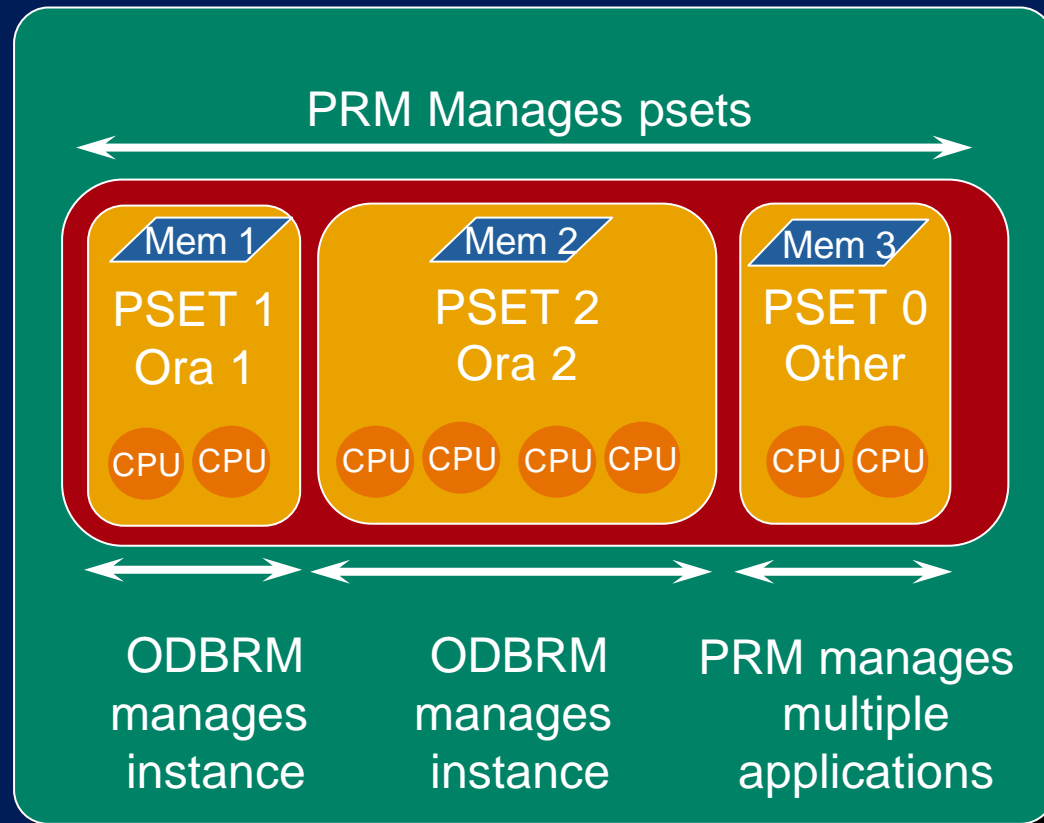


- Resource partitions within a single OS image



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

Integrated hierarchical management 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 9000 Server
 Processor set defined by PRM
 Memory allocated by PRM

Workload Management



Target Problem



- Handling Peak Demand for Critical Applications



Traditional Approach



- **Overprovisioning**

- Lots of dedicated Unix servers
- Excess capacity on each

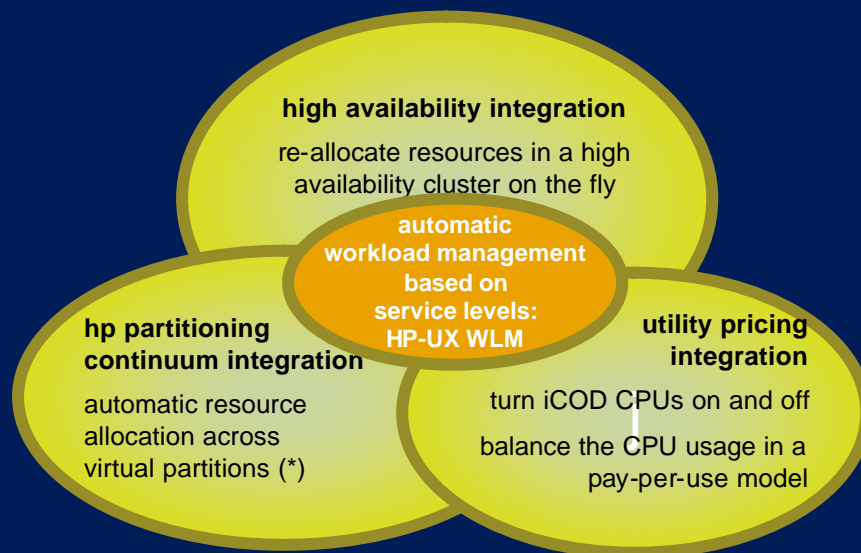
- **Drawbacks**

- Cost of underutilised capacity
- Difficult to manage many systems

workload manager



- goal-based resource management – automatic allocation of CPU resources based on set service level objectives both within an OS image and across vPars
- Predictable response times for mission-critical applications
- process resource manager included



out-of-the box toolkits

- Oracle database toolkit to collect
- Oracle performance metrics
- Apache toolkit
- SAS software toolkit

increased optimization capabilities through integration with

- servicecontrol manager
- virtual partitions
- MC/serviceguard
- iCOD and pay-per-use
- Application Response Measurement (ARM)

- Direct measurement of the performance of the workload
 - Response Time
 - Throughput
- Measurement of load on application
 - Number of users/processes
 - Queue length
- Resource Usage
 - CPU entitlement based on utilisation of current entitlement

Examples of Service Level Objectives (SLOs)

application a

application b

application c

response time SLO

transactions will complete
in less than 2 seconds.

priority 1

response time SLO

transaction will complete
in less than 3 seconds

priority 2

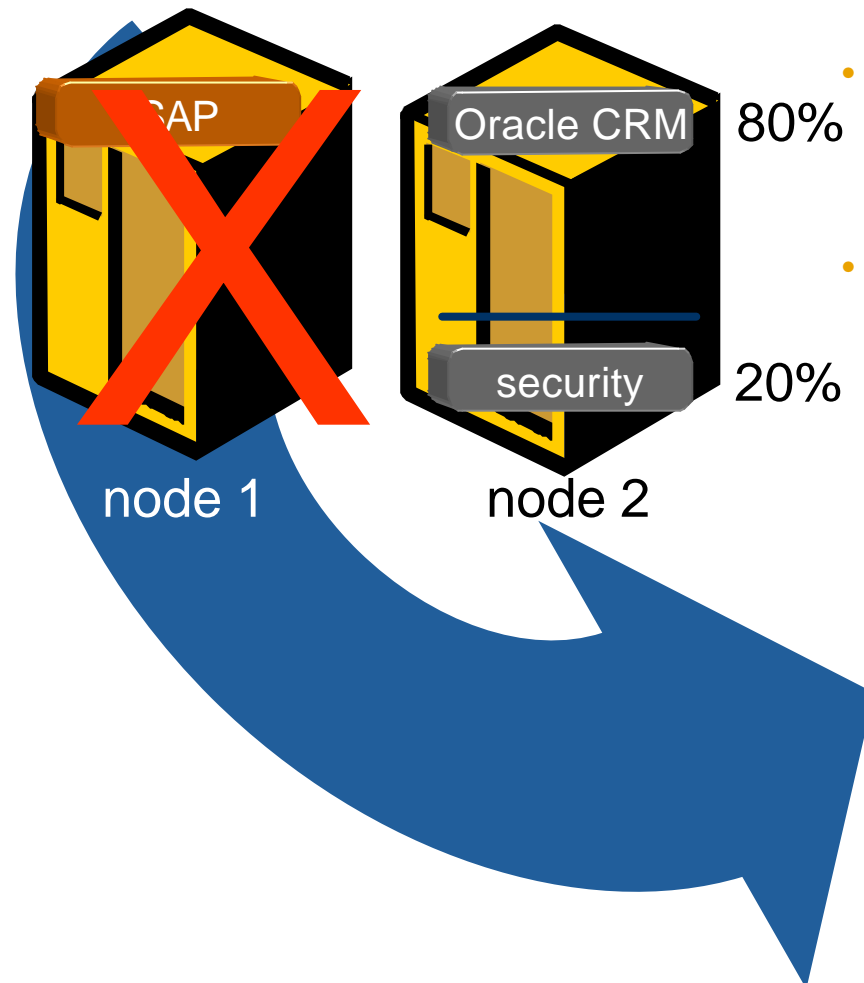
job duration SLO

batch job will finish in
less than 1 hour.

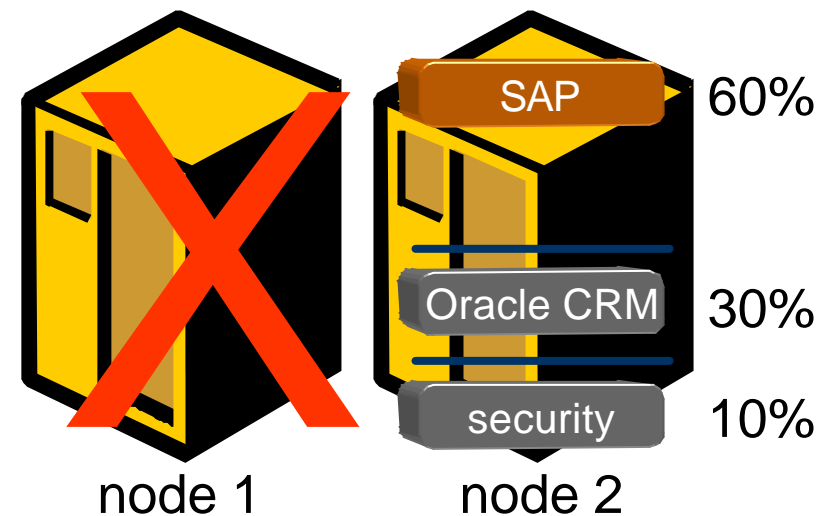
priority 3

HP-UX WLM automatically reconfigures CPU
resources to satisfy SLOs in priority order

Serviceguard and HP workload manager



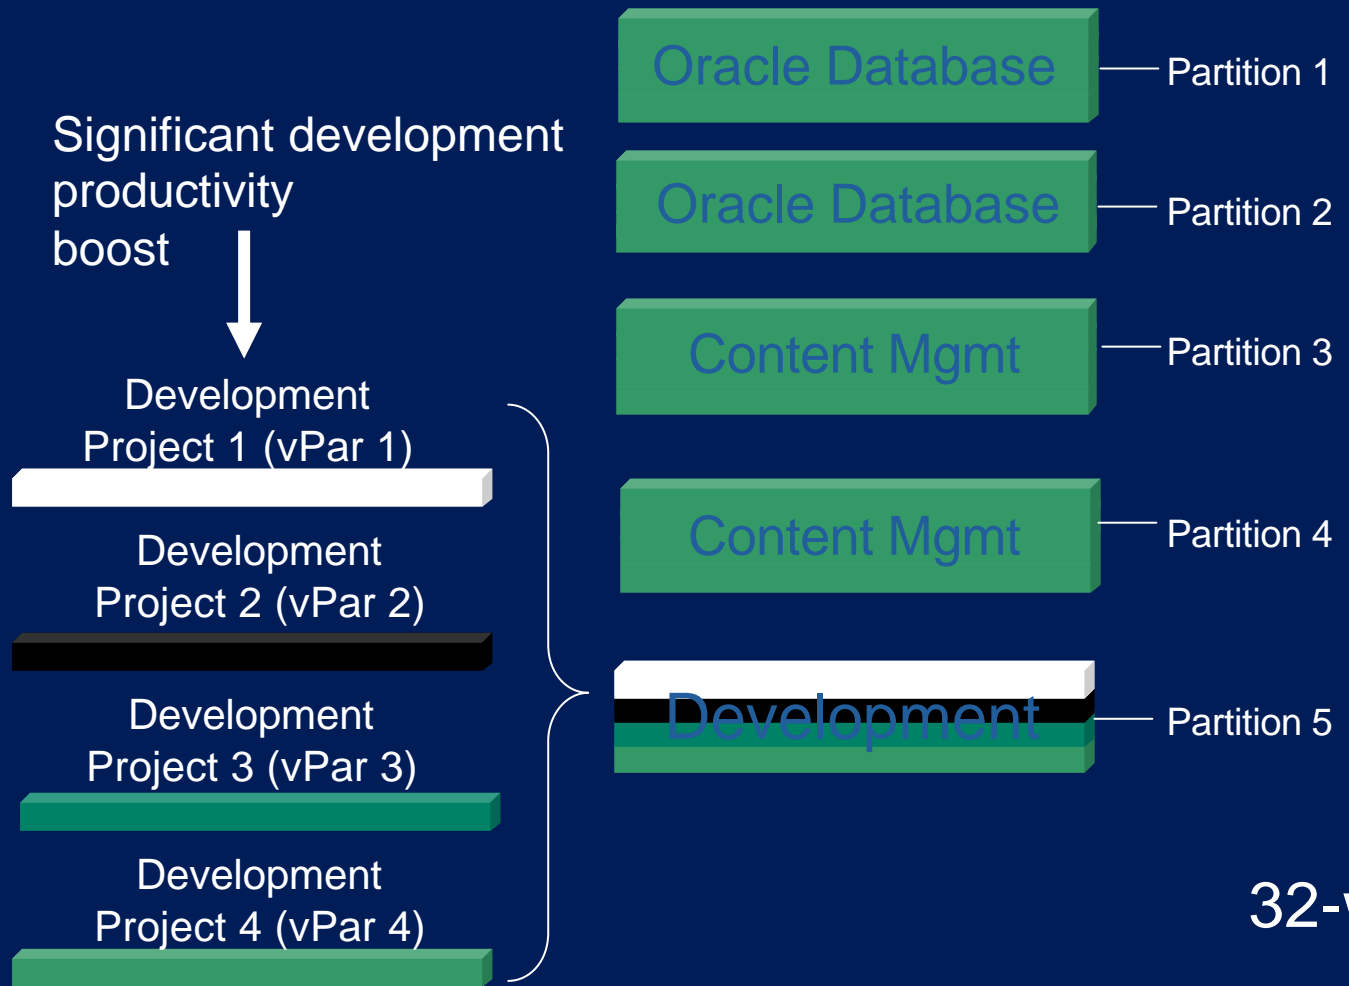
- **ServiceGuard**
 - analyses new infrastructure configuration
 - automatically moves package(s)
- **Workload Manager**
 - automatically aware of new package
 - re-priorities packages & resources to maintain SLA
- **iCOD**
 - communicates with WLM
 - instantly activates additional processors to fulfill SLA commitments



Customer Examples



partition example: vpartitions to isolate development environments



32-way Superdome

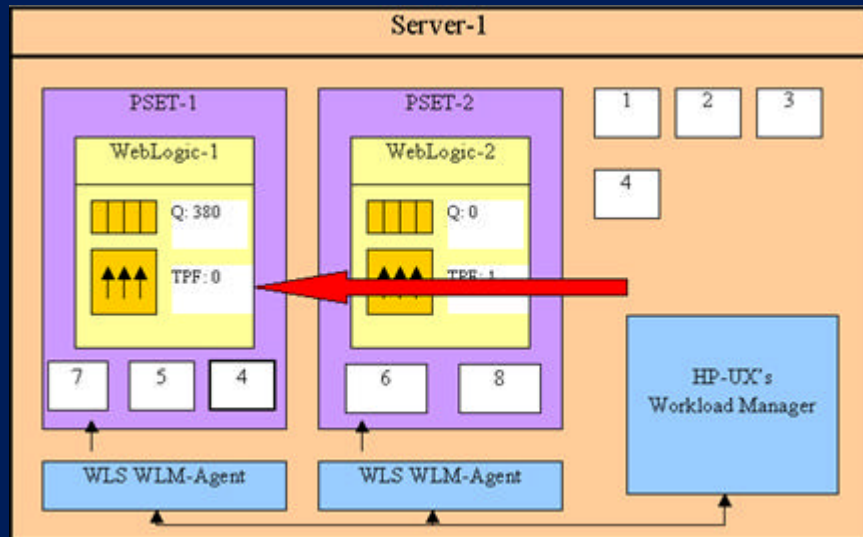
- *“HP-UX Workload Manager helped us to maximize system utilization and reduce our capacity management costs and also helped us allocate resources to multiple Oracle instances. In addition, the integration of HP-UX Workload Manager with MC/ServiceGuard helped us to prioritize production database instances on failover.”*

Tommy Simmons
*UNIX Team Leader
Wachovia Bank*

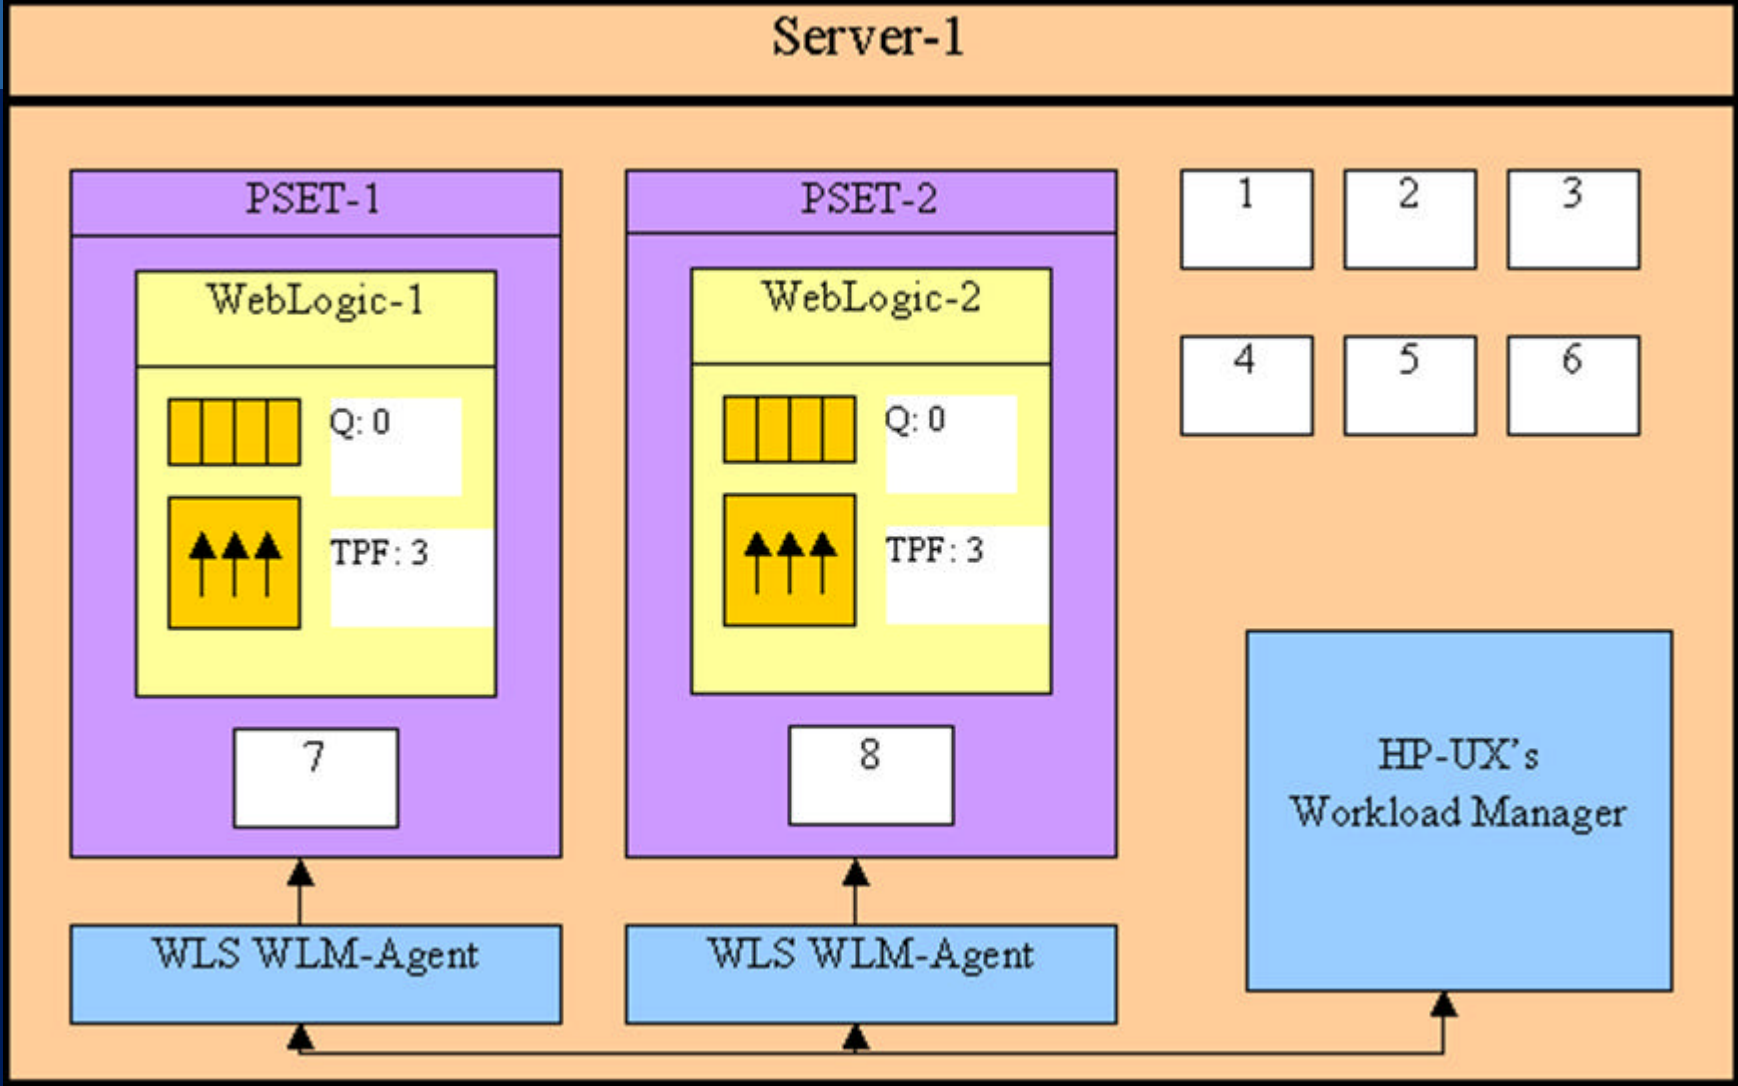
Consolidating Oracle and Weblogic on Superdome using HP-UX WLM



- Large Financial Services Company
 - 29,000 employees in 40 countries
 - 82-year history
 - Earned nearly \$1.8 billion in 2001
- Processor Sets provide the optimal performance and throughput for Weblogic-based applications
- Current queue length and the number of idle threads in the associated thread pool are used as performance metrics for WLM
- Workload Manager will dynamically resize processor sets for optimal performance



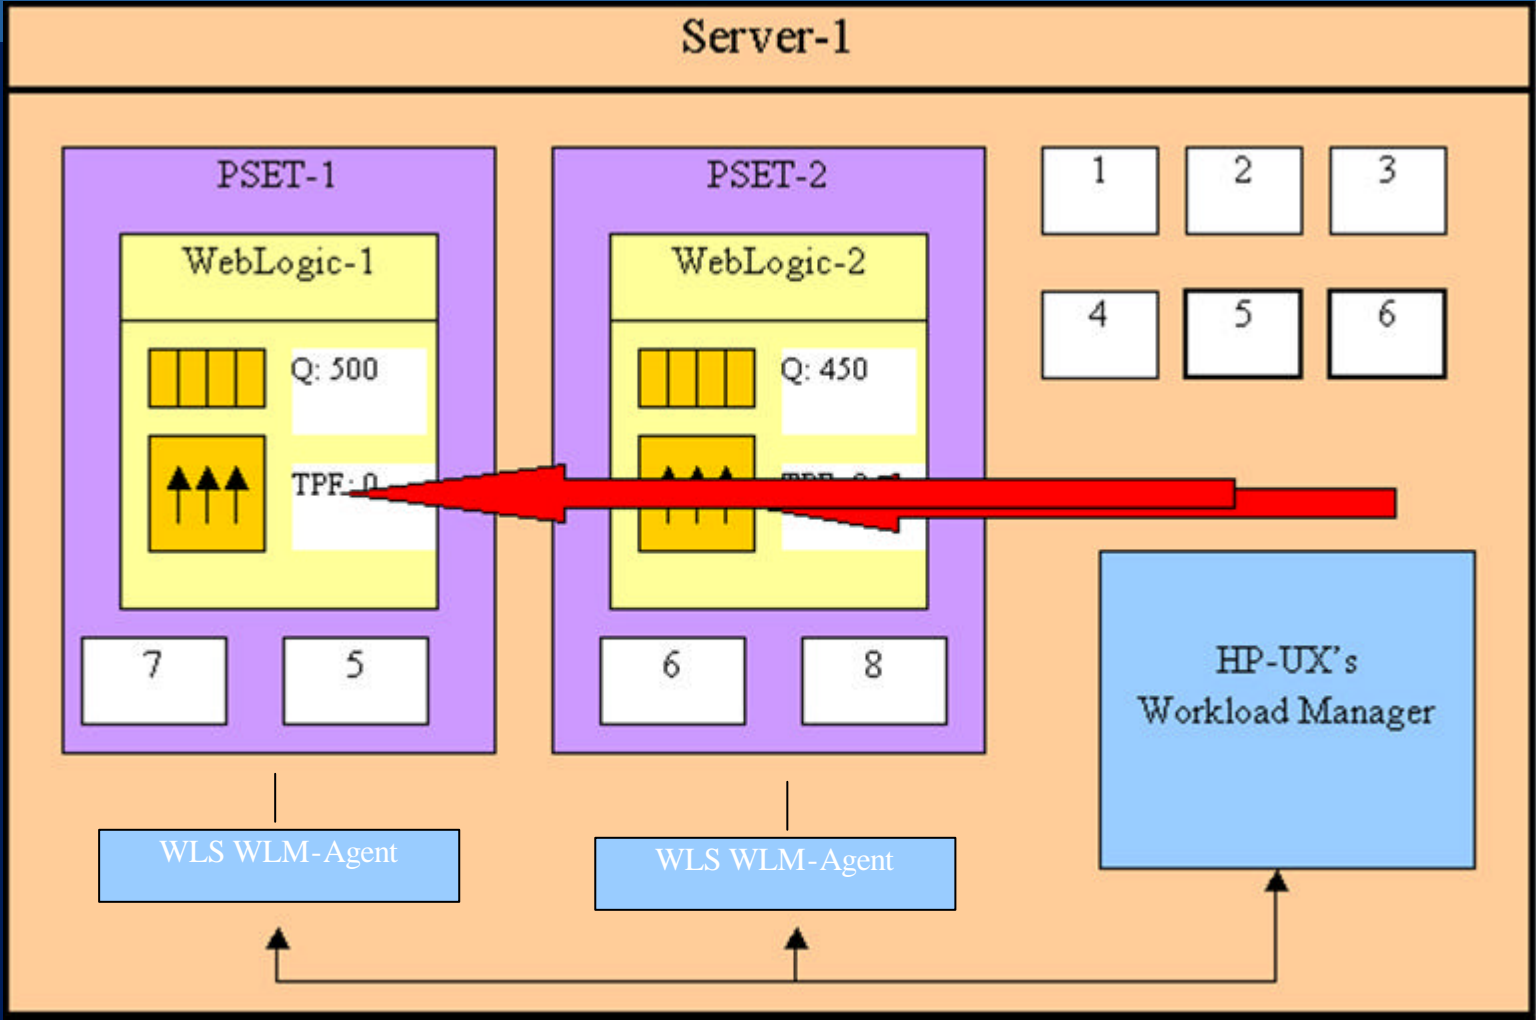
Weblogic / WLM integration



Key:

- CPU: 
- WebLogic Queue: 
- Thread Pool Free: 

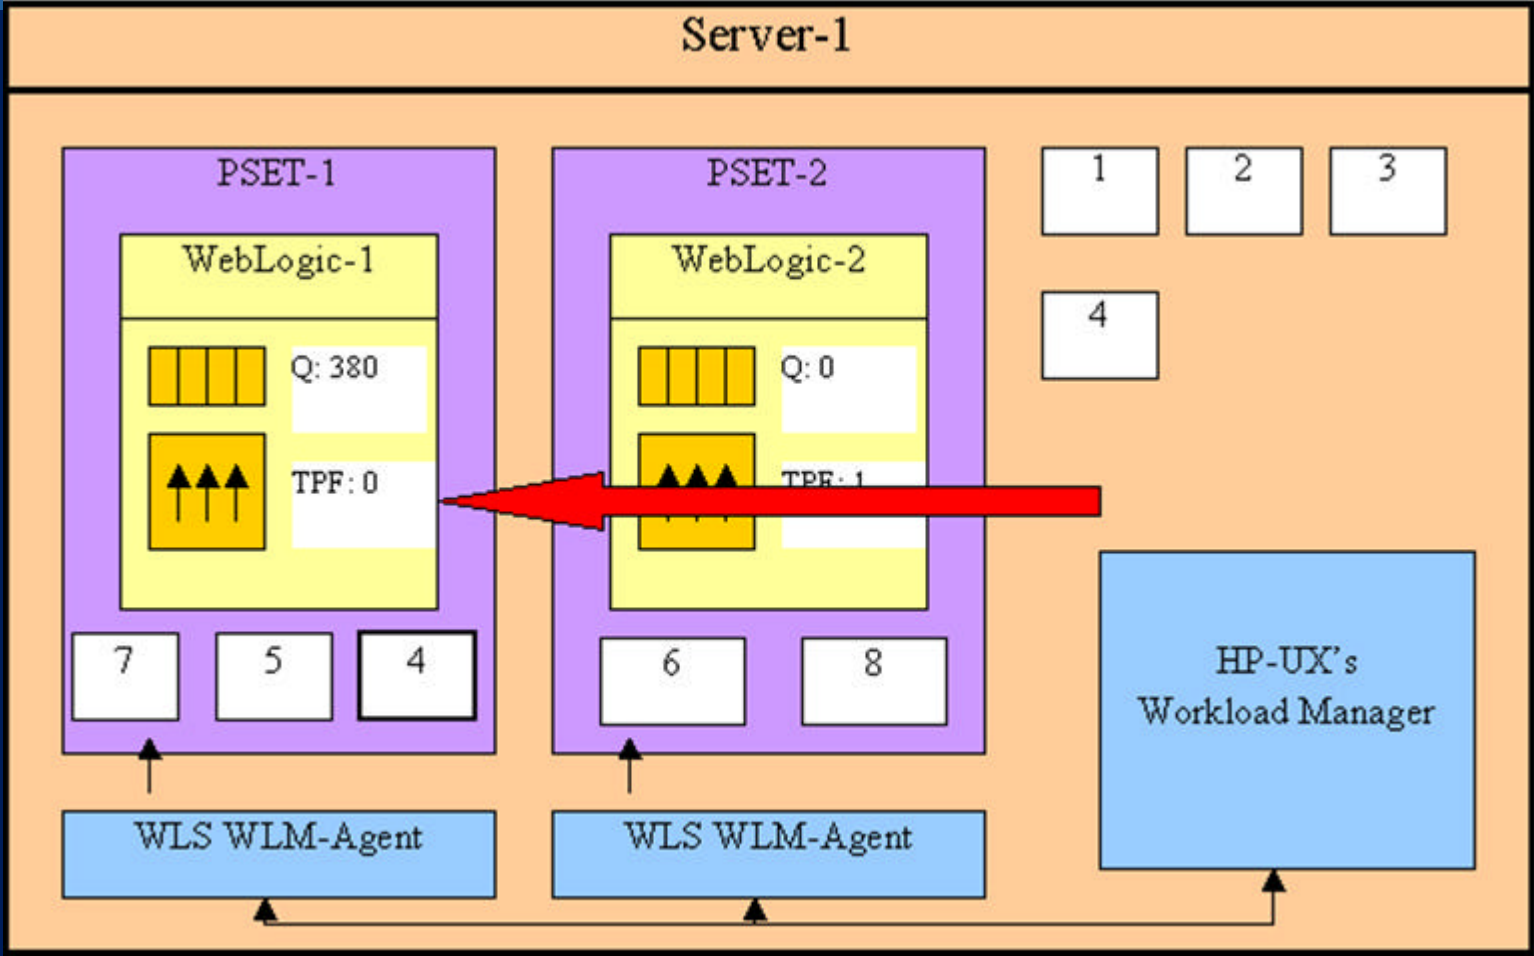
Weblogic / WLM integration



Key:

- CPU: 
- WebLogic Queue: 
- Thread Pool Free: 

Weblogic / WLM integration



Key:

- CPU:
- WebLogic Queue:
- Thread Pool Free:

Demonstration





i n v e n t