



HP Partitioning Continuum

hp world #432

Hewlett-Packard
Shefali Chinni

Pruneridge Ave
Cupertino, CA 95014

Phone: 408-447-2466

Fax: 408-447-4594

E mail: shefali_chinni@hp.com

HP Partitioning Continuum for Always On

dynamic and optimized
usage of compute
power with uptime



agenda

why is partitioning important?

definition of partitioning

HP Partitioning Continuum
for Always On

new: HP Virtual Partitions

workload management tools
(PRM/ HP-UX WLM)

HP's competitive advantage

summary

resources

Why is partitioning important?

Pressure to offer service level guarantee at reasonable costs

Under utilization of servers

Address high fluctuation of Web and App traffic

Flexibility with privacy and high availability

HP Partitioning Customer Benefits

Meet service level agreements with best return-on-investment

80-90% + Utilization of compute power

Fast and dynamic implementation of changing requirements

"right" level of application isolation with uptime




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.

HP Partitioning Continuum for Always On

HyperPlex
hard partitions with multiple nodes



nPartitions
hard partitions within a node

OS image with HW isolation

OS image with HW isolation

OS image with HW isolation

Virtual Partitions
within a hard partition

hard partition

OS image with SW isolation

OS image with SW isolation

OS image with SW isolation

PRM with psets
resource partitions within a single OS image

1 OS image

Application 1
with guaranteed compute resources

Application 2
with guaranteed compute resources

Application n
with guaranteed compute resources

Based on CPUs or percentages



hp-ux wlm (workload manager)

-automatic goal-based resource allocation via set SLOs



Isolation
highest degree of separation

Flexibility
highest degree of dynamic capabilities

hp partitioning continuum for always-on 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

**hp
hyperplex**

nPartitions

**Virtual
partitions**

PRM
(Process Resource
Manager)

- complete hardware and software isolation
- node granularity
- multiple OS images

- hardware isolation per cell
- complete software isolation
- cell granularity
- multiple OS images

- complete software isolation
- CPU granularity
- dynamic CPU migration
- multiple OS images

- dynamic resource allocation
- share (%) granularity
- 1 OS image

hp-ux wlm (workload manager)


- automatic goal-based resource allocation via set SLOs

Isolation
highest degree of separation

Flexibility
highest degree of dynamic capabilities

hp 9000 enterprise server product line powering the HyperPlex architecture

HyperPlex



N-Class



L-Class



A-Class



SuperDome



HyperPlex nodes

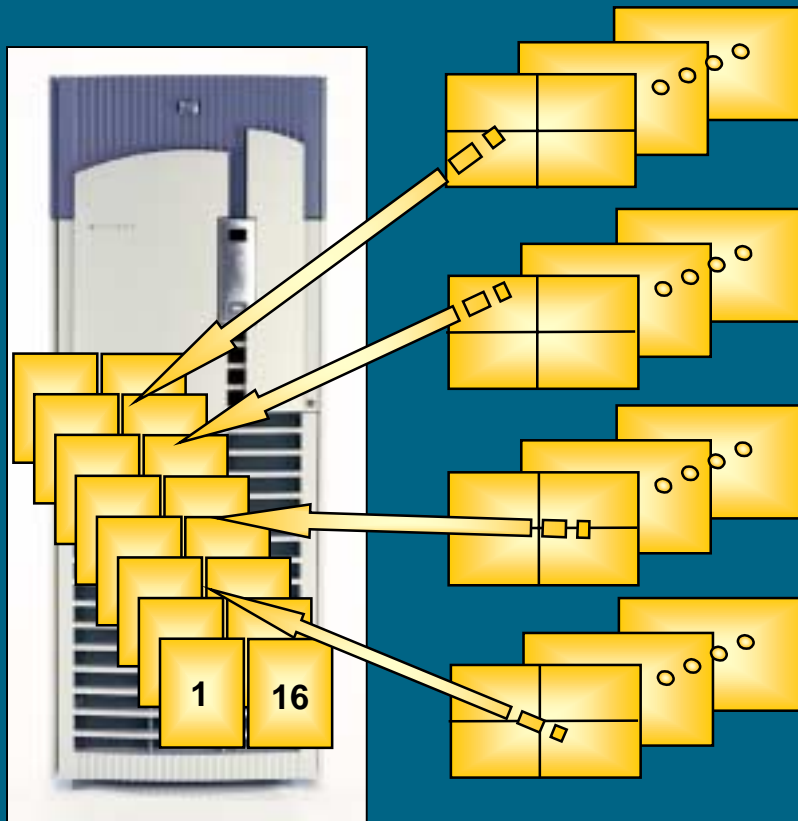
HP Confidential

R-, D-, K- and V-Class on CPL until otherwise noted

New!

nPartitions for SuperDome

Multiple applications
on the same server
with hardware isolation



Increased system utilization

- partitioning SuperDome 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 and software isolation across nPartitions
- MC/ServiceGuard support (within SuperDome or to another HP 9000 server)

HP Future

HP Virtual Partitions:

Multiple applications on the same server with software isolation

new!

core functionality for free!

Dept. A App 1	Dept. A App 1'	Dept. B App 2	Dept. B App 3
HP-UX Revision A.1	HP-UX Revision A.2	HP-UX Revision B.3	HP-UX Revision B.3



Increased system utilization

- partitioning a single physical server or hard partition into multiple virtual partitions for L-Class, N-Class, and Superdome

Increased Flexibility

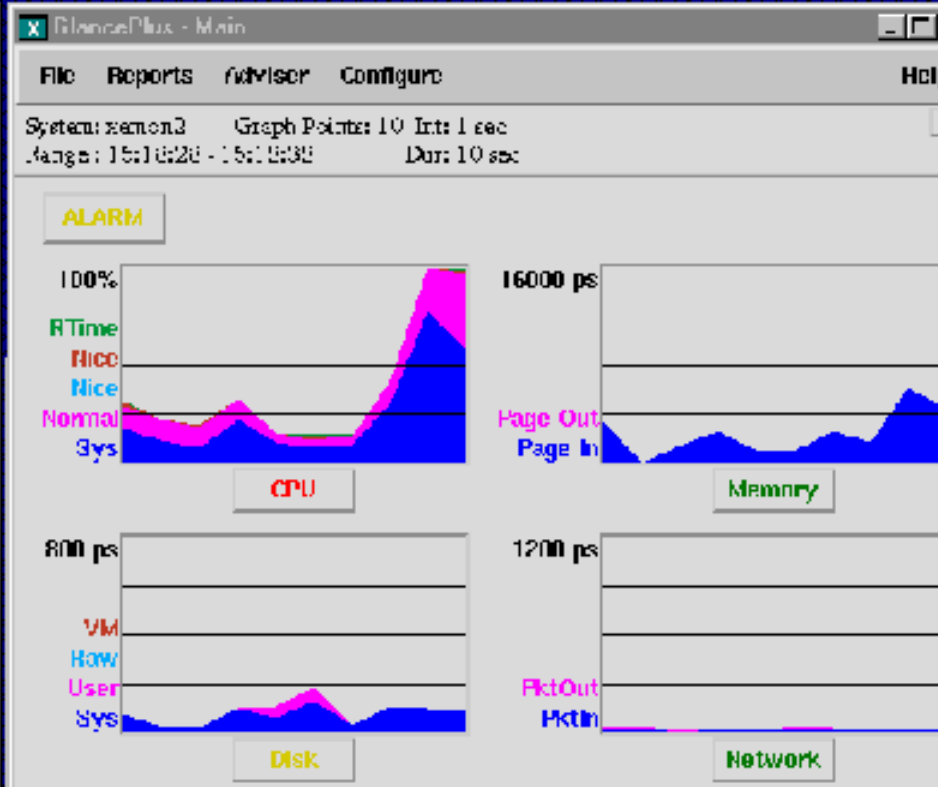
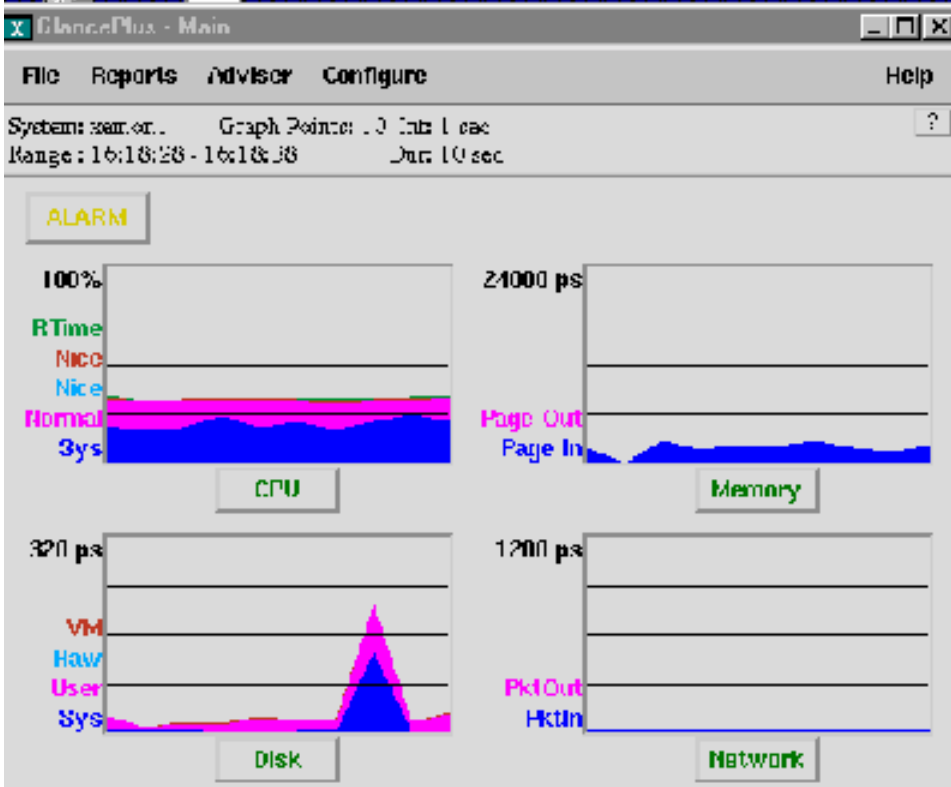
- multiple independent instances of HP-UX
- dynamic CPU migration across virtual partitions

Increased Isolation

- application isolation across virtual partitions
- OS isolation
- individual reconfiguration and reboot




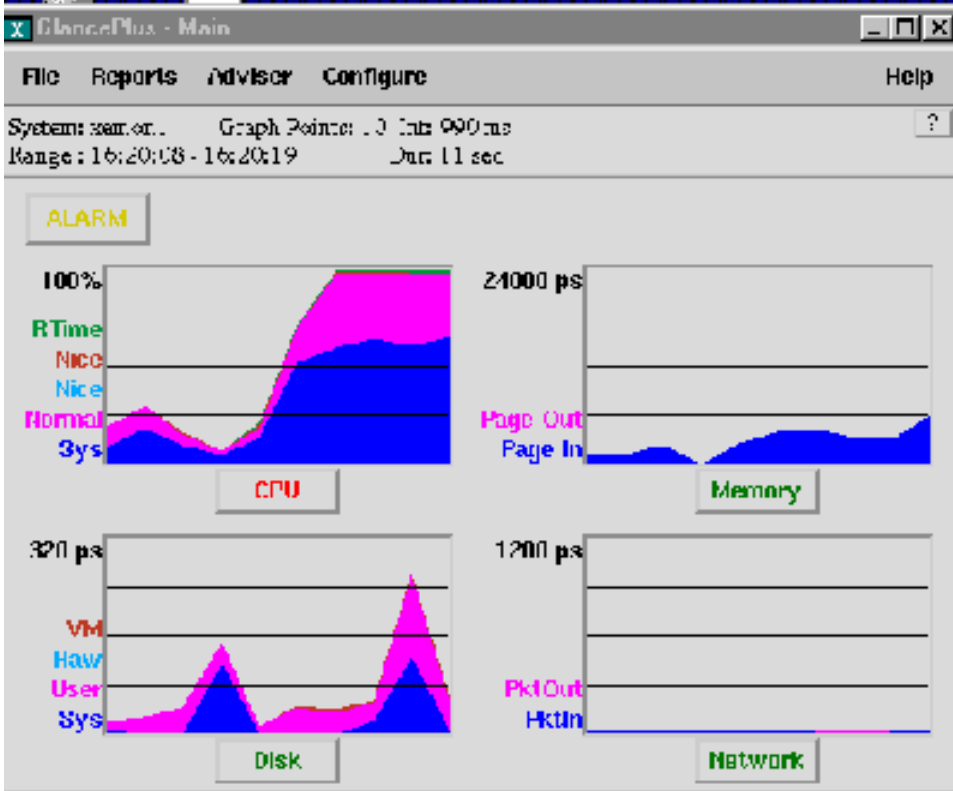
```
warp9
SDET result: 3109.3 scripts/sec
SDET result: 2443.8 scripts/sec
SDET result: 4933.6 scripts/sec
SDET result: 6343.6 scripts/sec
SDET result: 4413.8 scripts/sec
```

```
warp9
SDET result: 5530.5 scripts/sec
SDET result: 6216.9 scripts/sec
SDET result: 4838.3 scripts/sec
SDET result: 2437.6 scripts/sec
SDET result: 3127.9 scripts/sec
```



```
warp9
SDET result: 4933.6 scripts/sec
SDET result: 6343.6 scripts/sec
SDET result: 4413.8 scripts/sec
SDET result: 2455.8 scripts/sec
SDET result: 2454.2 scripts/sec
```

```
warp9
SDET result: 4838.3 scripts/sec
SDET result: 2437.6 scripts/sec
SDET result: 3127.9 scripts/sec
SDET result: 6479.2 scripts/sec
SDET result: 6389.4 scripts/sec
```



HP Virtual Partitions Key Features

- ⇒ core functionality bundled for free with every HP-UX 11i release
- ⇒ support of multiple HP-UX instances (HP-UX 11i and later)
- ⇒ different virtual partitions can run different versions of HP-UX
- ⇒ support of HP 9000 L-Class, N-Class, Superdome (including nPartition)
- ⇒ single CPU granularity (virtual partition may contain single CPU)
 - ⇒ L-Class - recommended up to 2 virtual partitions (max. 4)
 - ⇒ N-Class - recommended up to 4 virtual partitions (max. 8)
 - ⇒ Superdome - recommended up to 32 virtual partitions (max. 64)
- ⇒ dynamic CPU migration across virtual partitions
- ⇒ software fault isolation (application and OS isolation)
- ⇒ Individual reconfiguration and reboot, e.g. for rolling upgrades (virtual partitions don't affect each other)
- ⇒ command line interface (in future via GUI)
- ⇒ single toggle console (in future consolidated console)
- ⇒ compatibility with PRM, HP-UX WLM, ServiceControl Manager and MC/ServiceGuard

Shipment for L and N-Class in 3Q01(Superdome after L &N)

HP-UX processor sets (psets) overview

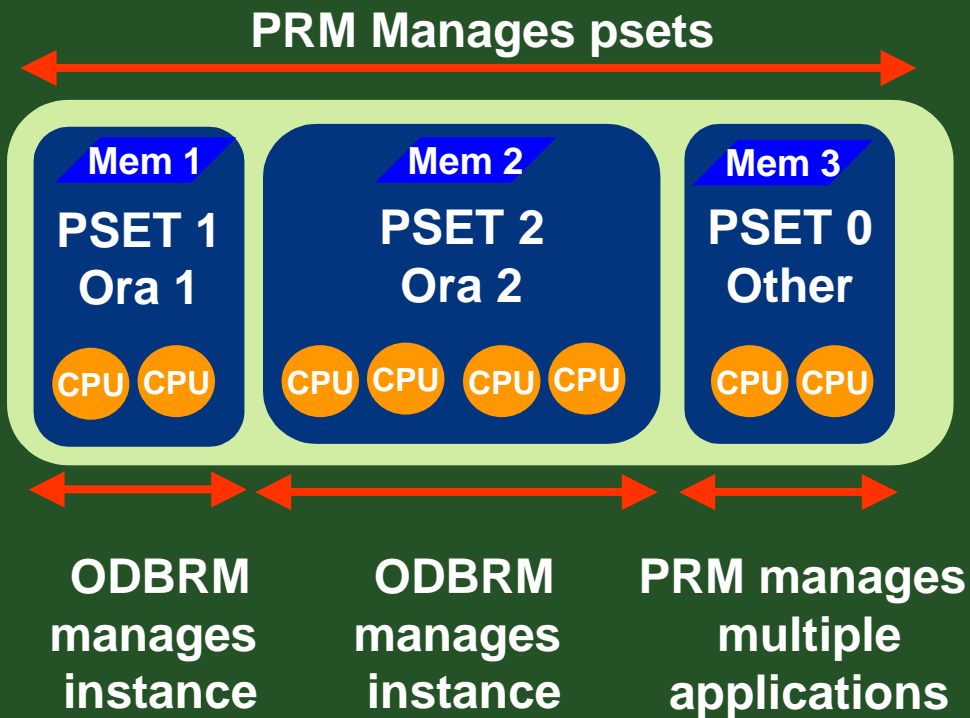
- represents a group of processors in the system
- represents a scheduling allocation domain
- provides mechanism for CPU resource management
- provides CPU resource isolation for applications and users
- does not provide fault isolation

the system may be configured into more than one processor set

HP-UX processor sets (psets) features

- dynamic creation, deletion, and reconfiguration of psets
- dynamic migration of threads and processes across psets
- ownership and access permissions for psets
- attributes to control psets behavior under different conditions
- processors are assigned to one pset at a time
- processes and threads have binding to one pset at a time
- system default pset for default users
- integration with PRM and gang scheduler
- can run PRM and ODBRM in different psets at the same time on a single system

integrated hierarchical management of Oracle on HP-UX



- HP 9000 Server
- Processor set defined by PRM
- Memory allocated by PRM

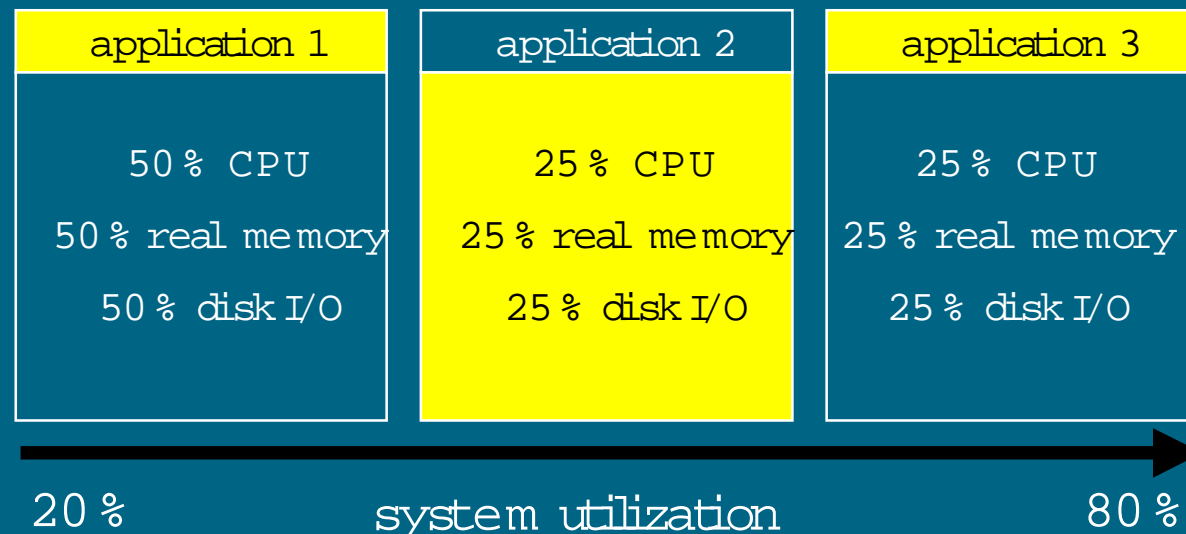
PRM cooperates with Oracle 9i DB Resource Manager to manage system resources

Result: HP 9000 is the ideal Oracle consolidation platform

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

PRM includes:

- PRM analyze to collect system accounting data for charge back and tuning



**Sold
>13,000
licenses**

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

workload management news

new features of PRM as of june 2000

- hierarchies
 - allows allocation of system resources for subgroups
- in addition to percentage-based allocation PRM also supports shares now
- in-kernel memory (supported with HP-UX 11i)
 - improves memory allocation
- single-point administration, Java-based GUI
 - runs standalone or within a browser: HP-UX 11.x, NT 4.0
 - supports IE 5.0 or Netscape 4.7
 - allows configurations of multiple OS images from a single console

(*) PRM release 1.08 for HP-UX 11 only

HP-UX workload manager (WLM)

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.

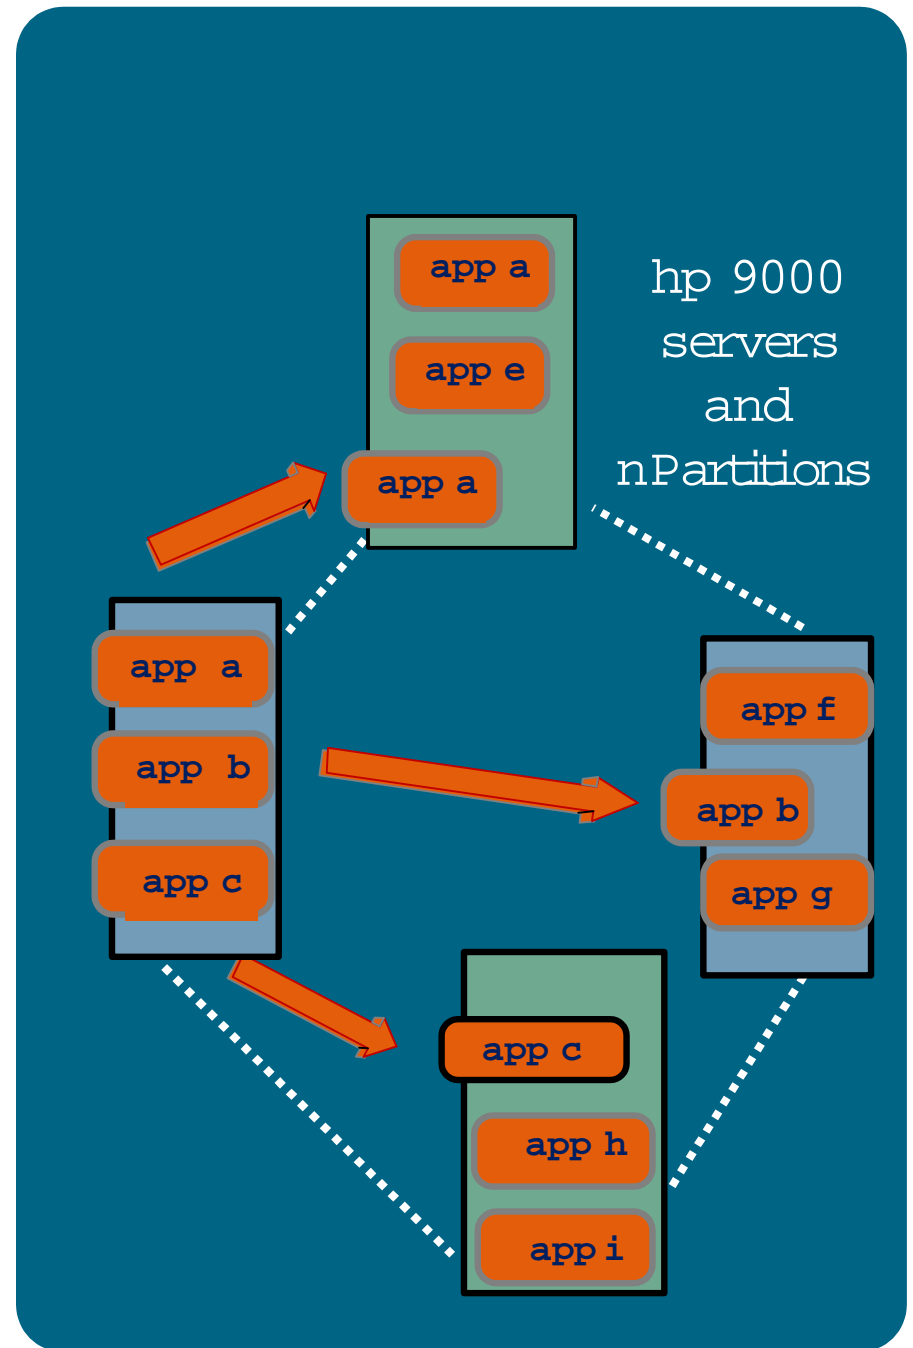
priority 3

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

dynamic and automatic application rehosting and reallocation

with mc/serviceguard
and hp-ux wlm

- minimize planned service interruptions
- hardware/software upgrades or maintenance
- goal-based resource allocation with hp-ux wlm based on set SLOs





hp-ux workload manager (WLM) - new release: hp-ux WLM 1.1 Features



- Application Response Measurement (ARM) Integration Toolkit (ARM-IT)
 - Enables automatic management of ARM'ed applications with WLM
- Improved Data Collection Capabilities
 - Command line interfaces enable data collection with script languages
- MC/ServiceGuard Integration
 - Workload management with a single configuration file throughout a MC/ServiceGuard cluster
- Selective Capping
 - Allows user-selected workloads to share CPU entitlements when not in use while limiting others to their entitlements.
- Latest version of PRM (release 1.08) included
 - hierarchies, in-kernel memory, and shares

growing within a system or partition



application A

CPU

CPU

CPU

CPU

application B

CPU

CPU

CPU

CPU



application A

CPU

CPU

CPU

CPU

CPU

CPU

CPU

CPU

application B

CPU

iCOD CPUs

CPU

active CPUs

CPU

iCOD CPUs turned on

iCOD notification:
WLM turns the iCOD CPUs on in order to
manage the workload

HP-UX Workload Manager (WLM)

New Release:

HP-UX WLM 1.2 Features



Shipment
June 01

- Oracle Tool kit for hp-ux WLM
 - collection of the performance metrics for Oracle made easy
- Integration with ServiceControl Manager
 - launch WLM from the SCM GUI for rapid deployment
- iCOD notification
 - to alert the system administrator for more iCOD reserves needed to meet the SLOs

managing hp partitioning continuum

the integrated power of
servicecontrol and openview



servicecontrol manager

**single-point multi partition
management for fast
deployment and
consistency**

partition manager

**setup partitions
map partition architecture**

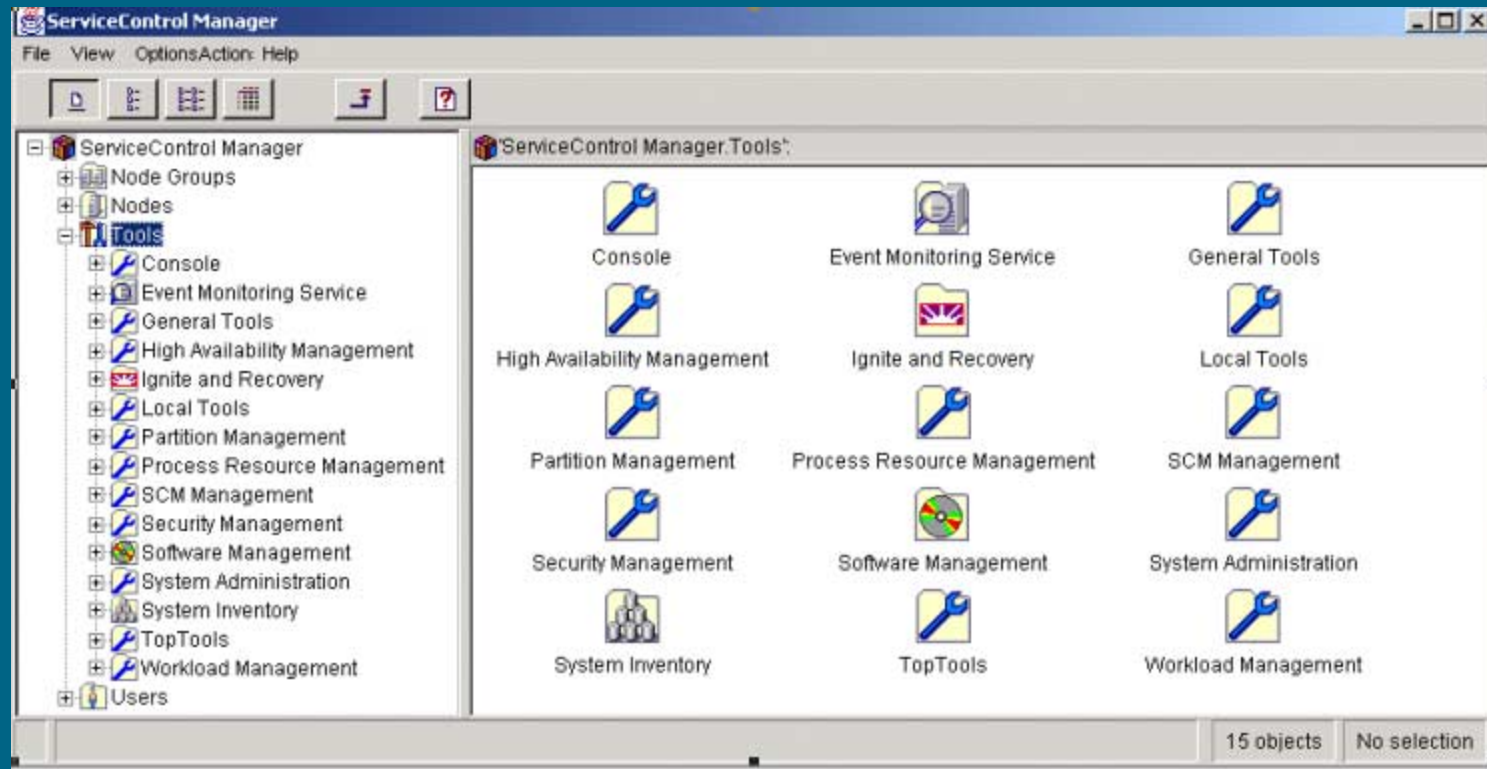
openview glanceplus pak

**performance monitoring of
each partition**

**openview vp operations
monitor events on each
partition**

deploy

servicecontrol manager tools - news manage multiple partitions



All hp-ux manageability tools are integrated - in addition to toptools for device management and Linux agent support to manage Linux-based server nodes

TopTools (Aug 2001)

Process Resource Manager (Sep 2001)

Serviceguard Manager (Sep 2001)

Central Web Console (Sep 2001)

Partition Manager (Sep 2001)

SCM Management (Oct 01)

deploy

servicecontrol manager new release: new scm 2.0 features

shipment
june 2001

increased scalability - 16x

- up to 1024 managed nodes per managed cluster

increased security

- public/private key authentication
- "on the wire" encryption by using HP IPSEC product

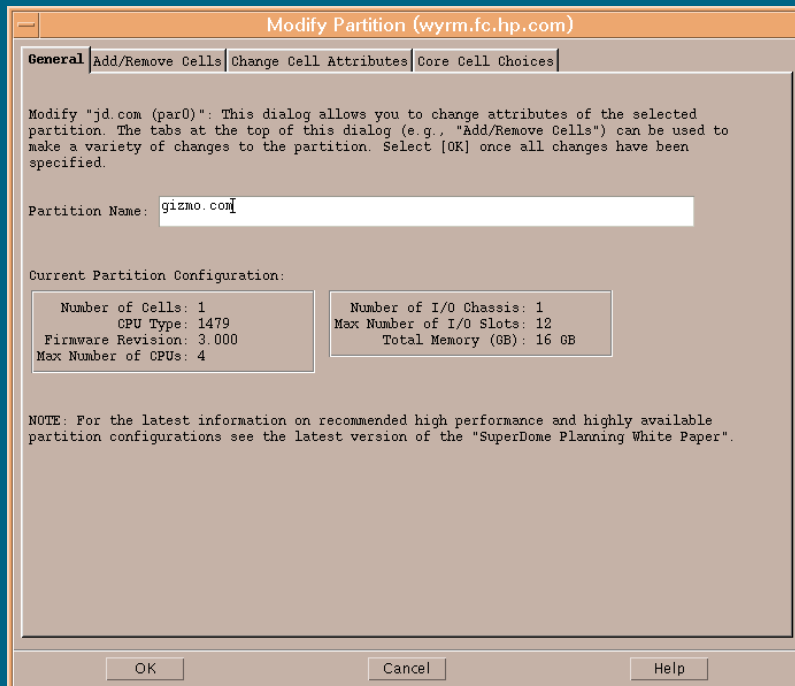
increased productivity

- toptools event log viewer
- node group authorization

investment protection

- direct upgrade path from scm 1.0 and 1.1 to 2.0

partition manager



Note: GUI for nPartition, GUI for virtual partitions planned

- Create & modify partitions
- Display a complete hardware inventory
- Display status of key complex components
- Check for problem or unusual complex conditions
- Manage power to cells & I/O chassis
- Turn on/off attention indicators for cells, I/O chassis, I/O cards & cabinets

Keystone Systems Consolidation Sample Configuration



nPartition 1
for test environment

nPartition 2
for production environment

Virtual partition 1
for customer 1
with 1 or more
applications

Virtual partition 2
for customer 2
with 1 or more
applications

Test and
production
environment
with hardware
isolation

2 production
environments
with complete
software
isolation for
different
customers




Managing Partitions



nPartition 1
for test environment

nPartition 2
for production
environment



Glance Plus Pak
performance monitoring of partitions

Parmgr

To create and
configure partitions

ServiceControl
Manager

Manages partitions
for fast deployment
and consistency

hp partitioning continuum for always-on

THE competitive advantage

	HP-UX	Sun (Solaris)	IBM (AIX)
Resource partitioning	Yes prm with integrated psets (Sep 2001)	Yes psets not integrated with srm \$\$	Yes psets not available
Goal-based resource partitioning	Yes MC/Serviceguard and iCOD integration	No	No
Virtual partitions with S/W isolation	Yes	No	No
Hard partition within a node	Yes	Yes Less HA	No
Hard partition with multiple nodes	Yes	Yes No high speed connect	Yes \$\$

HP Partitioning Continuum for Always On - Summary

Point

- Broadest partitioning offering - Multiple OS images with H W & S W isolation and individual resource allocation
- Optimization through multiple applications on the same servers
- Dynamic and automatic resource allocation based on set SLOs

Customer Benefits

- "right" level of application isolation with uptime
- 80-90 % + Utilization of compute power
- Fast and dynamic implementation of changing requirements

Competitive Differentiation

- SUN Solaris and IBM AIX offer only subset - HP only UNIX vendor with virtual partitions
- SUN Solaris and IBM AIX with significant lower utilization
- SUN Solaris and IBM AIX have NO automatic resource allocation solution

Only with HP - meet service level agreements with best ROI

HP Partitioning Continuum for Always On

dynamic and optimized
usage of compute
power with uptime





backup

HP Partitioning Continuum for Always On

Decision Criteria

	HyperPlex	nPartitions	Virtual Partitions	PRM HP-UX WLM
Hard ware isolation need	SMP node	Cell	future	
Software isolation need	YES - multiple OS images	YES - multiple OS images	YES - multiple OS images	NO - 1 OS image
Granularity	SMP node	Cell	CPU	Share (%)
Dynamic Resizing		Yes (for PCI IO for cells - future)	Yes	Yes
	Isolation highest degree of separation		Flexibility highest degree of dynamic capabilities	

Customer Reference



SIEMENS

"Siemens Business Services will be enabled with HP Virtual Partitions to further optimize the system utilization by running multiple applications on the same HP 9000 server with complete software isolation. This new system consolidation capability will result in less servers to be managed in a more flexible way and ultimately lower total cost of ownership."

Helmut Kutzberger
UNIX System Management
Siemens Business Services

workload management success PRM customers



Samsung Securities Co.,Ltd.



QUELLE.

Willkommen im Land
der ungeahnten Möglichkeiten



Delta Air Lines



- sixth largest bank in USA
- more than 16 million customers who expect 24x7 access to account information.
- mission-critical solutions: 60% of the 330 hp 9000 v-class and k-class enterprise servers run mc/serviceguard

"First Union has had measurable success using HP's Process Resource Manager to manage multiple Oracle databases for traditional and web applications."

"The goal-based capabilities of Workload Manager will further enhance our e-service Center's ability to deliver the service levels our customers demand."

- Rob Young
First Union Corporation
December 1999