



i n v e n t



Technical Services



vPars Presentation



Asghar Riahi
Senior Technical Consultant
Hewlett-Packard
Pleasanton, California
Technical Services
Engineering Services West Region
Phone: 925-460-1627
Email: asghar.riahi@hp.com



Agenda

- What is Virtual Partitioning
- HP Partitioning Continuum for Always On
- Where HP vPars provides the most value
- Definition of Partitioning
- vPars w/ WLM
- What is vPars
- vPars Components Overview
- How vPars works
- Partitioning Plan
- vPars commands
- vPars Installation
- MC/SG and vPars
- vPars GUI vparmgr
- Summary



What is Virtual 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

- In traditional server environments, all CPUs within a server run a single OS instance and one or more applications. Application and OS failures may affect the entire system.
- With HP-UX Virtual Partitions, however, the software stack contains the vPar Monitor between the hardware/firmware level and OS, allowing each partition to support a different instance of HP-UX 11i (or later) with each instance capable of being at a different version and patch level.




HP Partitioning Continuum for Always On


HyperPlex
hard partitions with multiple nodes
(City w/ houses)



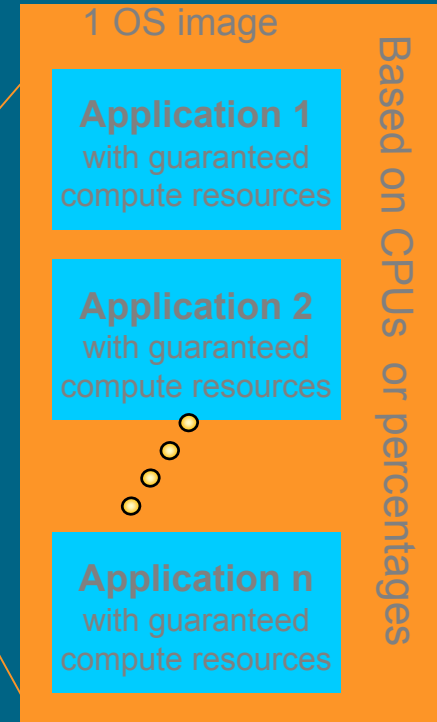
nPartitions
hard partitions within a node
(House w/ apartments)



Virtual Partitions within a hard partition
(Apartments with rooms)



Resource partitions within a single OS image
PRM with psets
(Rooms with closets)



hp-ux wlm (workload manager)

—automatic goal-based resource allocation via set SLOs

Isolation
highest degree of separation

Flexibility
highest degree of dynamic capabilities



Partitioning continuum
HP-UX partitioning offerings are all integrated.


They have been designed to complement each other to provide greater flexibility.

Hard Partitions with multiple nodes

Hard Partitions within a node

Virtual Partitions within a hard partition

PRM with psets resource partitions within an OS image




- complete hardware and software isolation
- node granularity
- multiple OS images
- PCI OLAR support



- hardware isolation per cell
- complete software isolation
- cell granularity
- multiple OS images
- PCI OLAR support




- software isolation
- multiple OS images
- dynamic CPU migration
- 1 CPU granularity
- available on low & mid-range servers
- resources not tied to physical configuration
- PCI OLAR support



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

WLM - workload manager): automatic goal-based resource allocation via set SLOs



Where HP vPars provides the most value

- ① **System/data center consolidation**
- ① **development/test environments**
- ① **increased system utilization**
- ① **varying workload requirements:**
 - **time of day: order entry during day, batch at night**
 - **time of month (payroll, end-of-month/end-of-year financials)**
 - **as particular needs require**
- ① **service provider (providing system resources to different users/applications)**
- ① **unique application tuning of O/S**
- ① **time zoning**



Definition of Partitioning

Partitions are physical or logical mechanisms which:

- Isolate operational environments within single or multiple servers
- Offer flexibility of dynamic resizing
- Protect applications from unrelated events that could cause disruption, interruption, or performance degradation



Hard partitions

Electrical isolation

Can perform hardware maintenance on one without affecting the others

Suitable for mixing production and R&D environments

Suitable for hosting multiple customers or entities

Soft/virtual partitions

Software isolation

Cannot perform hardware maintenance on one without affecting the others

Not suitable for mixing R&D and production environments

Not suitable for mixing customers or different financial entities

vPars w/ WLM: automatic goal-based resource allocation between vPars based on SLOs

Dynamic adjustment of resources by migration of CPUs between vpars when required

WLM

WLM Rules Sets

Business need defines the relative priority and response time for the various applications

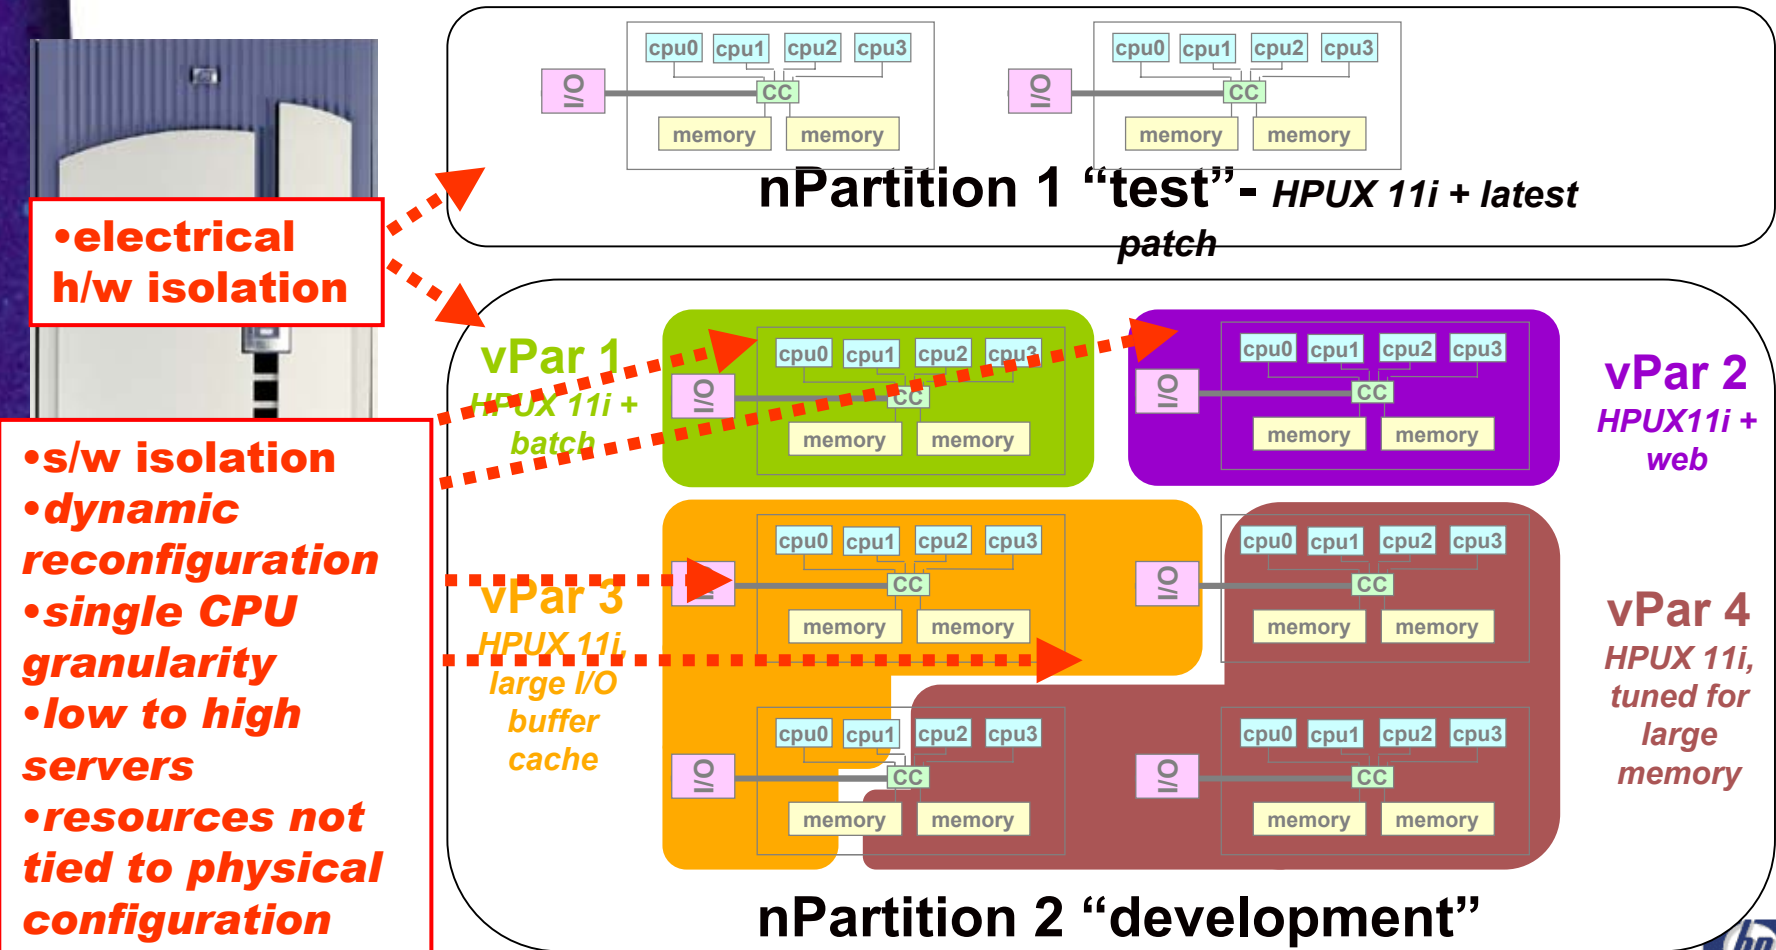
Dynamic feedback from vpars regarding load and response times

Dept. B App 2	Dept. B App 1	Dept. A App 1
HP-UX Rev Z	HP-UX Rev Y	HP-UX Rev X



HP Partitioning- vPars within nPars

- 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



A.02
What's new?



INCREASED FLEXIBILITY:

- ✓ Integration with nPartitions
- ✓ Integration with iCOD

SIMPLIFIED MANAGEMENT:

- ✓ Automatic, SLO-based workload management ACROSS vPars (WLM cross-vPar integration)
- ✓ Availability of new vparmgr GUI

EASIER DEPLOYMENT:

- ✓ Ignite/UX is now vPar-aware (special WINSTALL file not needed)

A.02.02 PLATFORM SUPPORT:

- ✓ rp5405, rp5470 (L3000), rp7400/ N400, rp7405, rp7410, rp8400, Superdome

**1st Unix
software
partitioning**

new!

HP Virtual Partitions:

Multiple applications
on the same server
with software isolation

**core functionality
for free!**

**New w/
HP-UX 11i**

Dept. A
App 1

HP-UX
Revisio
n A.1

Dept. A
App 1'

HP-UX
Revisio
n A.2

Dept. B
App 2

HP-UX
Revisio
n B.3

Dept. B
App 3

HP-UX
Revisio
n B.3

vPar Monitor

**Built on HP
resiliency**



**with performance
& availability**

increased system utilization

- up to 80-90%, since additional partitions can use unused parts of system

greater flexibility

- multiple independent OSs
- ❖ 1 CPU granularity per vPar
- ❖ dynamic movement of CPU power between vPars
- ❖ resources not tied to physical config. (like hard partitions)
- ❖ vPars on low to high-end

isolation

- of OS, applications, resources
- individual reconfig & reboot

Hard Partitions

- Designed to isolate applications from single point of failures
- Full hardware isolation
 - consists of multiple nodes
 - single node
- Can be divided into virtual partitions



HP-UX Virtual Partitions

HP-UX Virtual Partitions (vPars)

Partitioning Technology that provides applications and operating system isolations

Virtual Partitions can be created to run multiple instances of HP-UX 11i simultaneously on:

- A Single Server Node
- nPartitions

Each partition is assigned its own resources and hosts its own set of applications

- Multiple applications or multiple instances of same application may coexist on single system



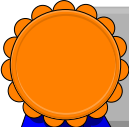

vPars Benefits

Benefits of using vPars

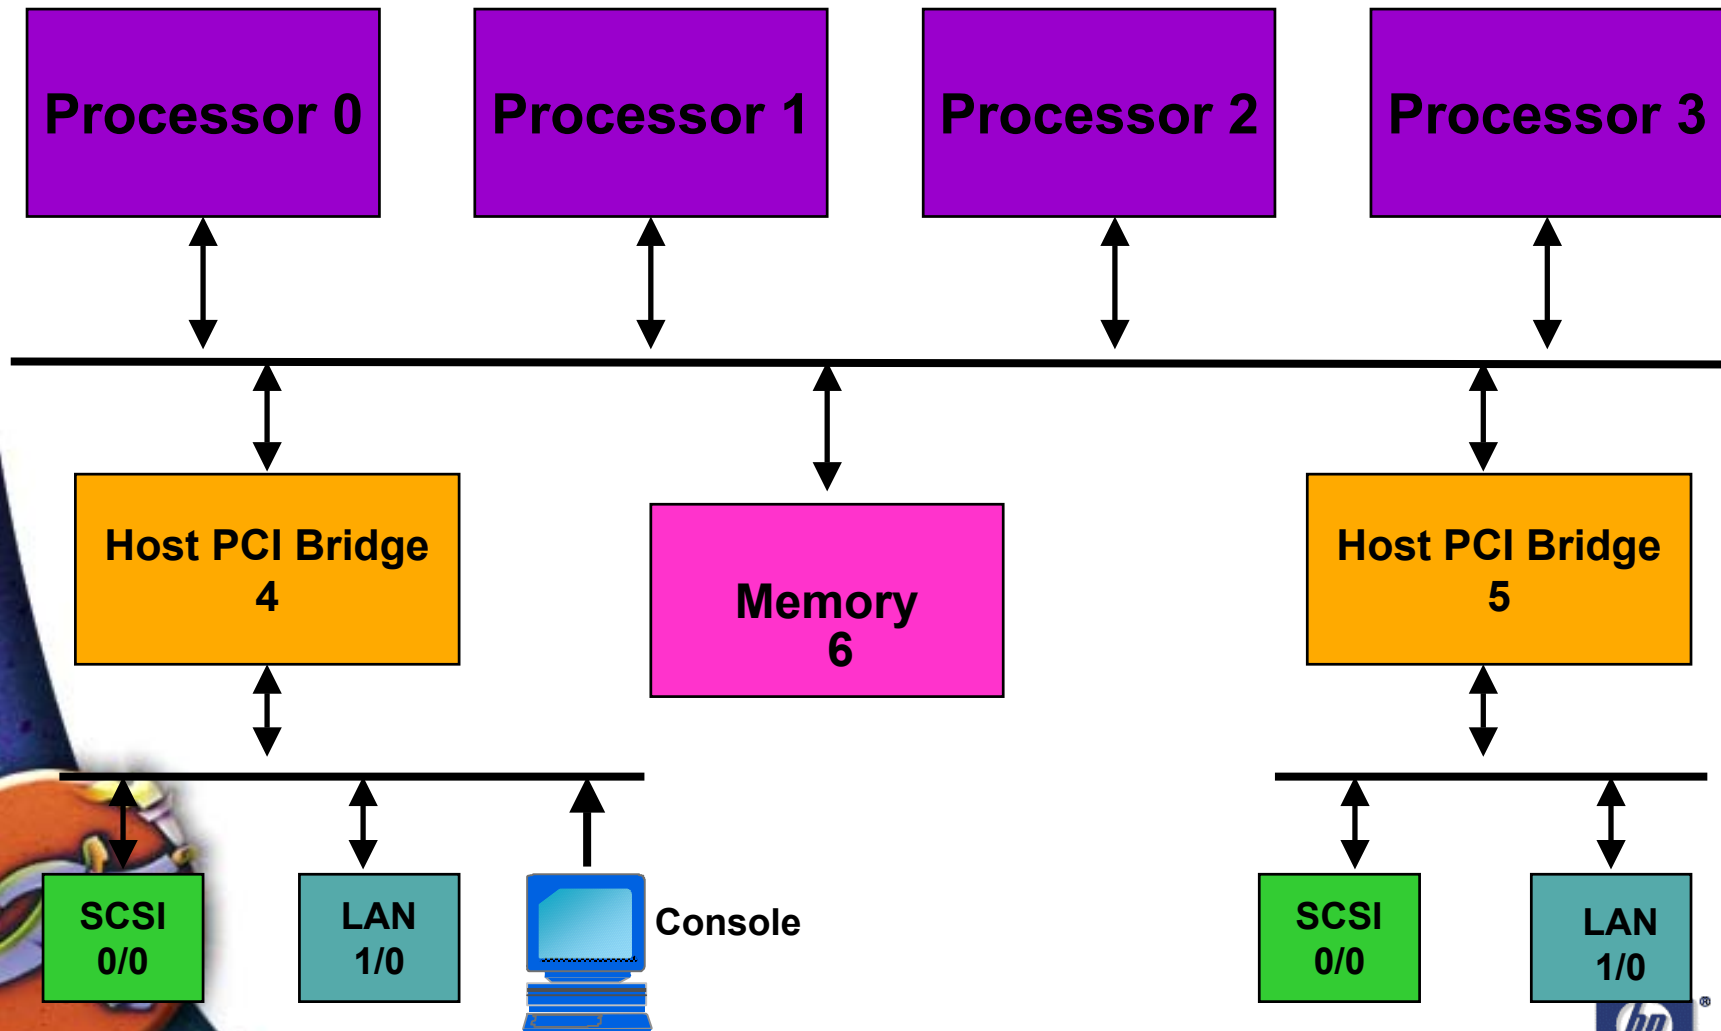
- Increased System Utilization
- Isolation of OS/OE & Applications faults
- Improved flexibility in configurations
- Improved system availability and scalability



The most comprehensive partitioning offering for the UNIX operating system

	 HP-UX 11i <small>Internet critical. E-services ready.</small>	Solaris™ 8	 AIX L v5.1
Hard partitions	Yes nPars	Yes	No
Soft partitions	Yes vPars	No	Yes
Integration with goal-based workload mgmt	Yes	No	No
Dynamic reconfig. of hard partitions	Yes	Yes	No
Dynamic reconfig. of soft partitions	Yes	No	No

System without vPars



System without vPars

➤ Software stack on system without vPars

Application

Application

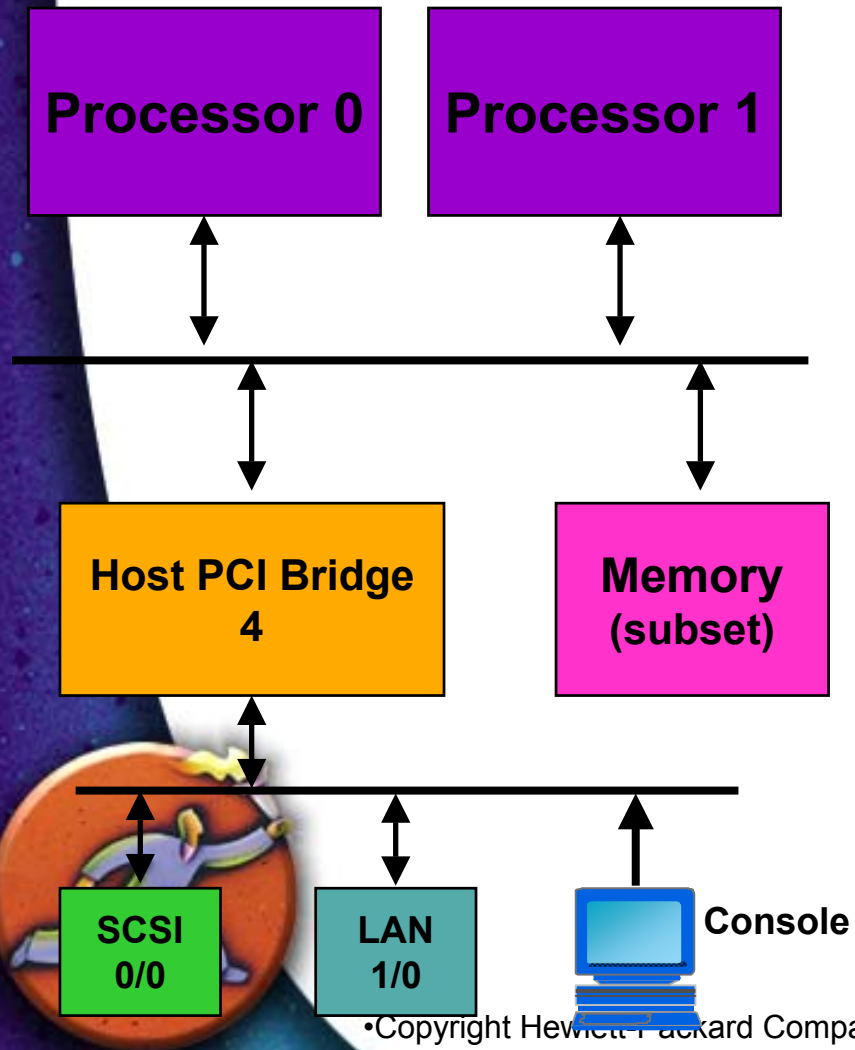
HP-UX 11i OE

Hardware / Firmware

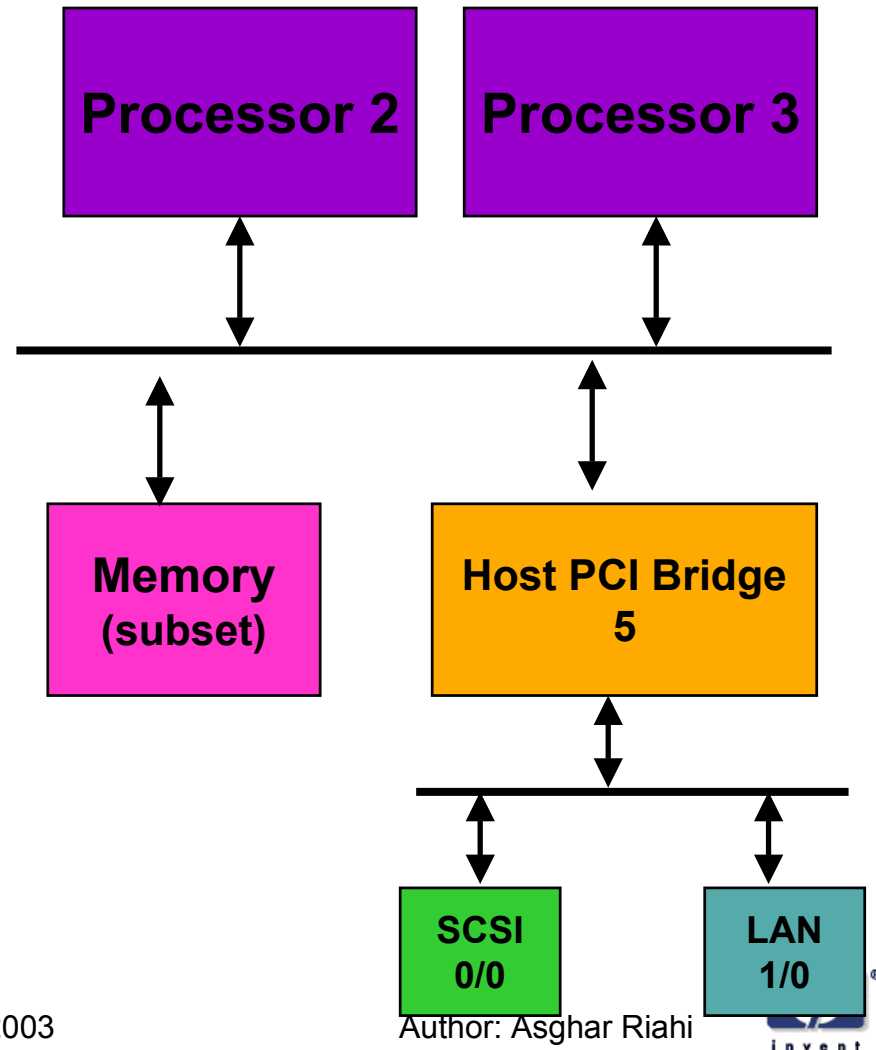


System with vPars

vPar 1



vPar 2



vPars Components Overview

- **vPars Monitor**
- **vPars Partition Database**



vPars Monitor

- **Manages the assignment of hardware resources to virtual partitions.**
- **Boots virtual partitions and their kernels.**
- **Emulates certain firmware calls**
 - *Creates the illusion that each vPar is a complete isolated HP-UX system*



vPars Partition Database

- **Contains partition configuration information**
- **Is located in vPars Monitor and a local copy is also stored on each virtual partition**
- **vPars Monitor uses the database to:**
 - *Track which virtual partition exist*
 - *Identify hardware resources and partition attributes that are associated with each partition*



vparstatus

- **SYNOPSIS**

```
vparstatus [-v | -M] [-p vp_name]... [-D  
db_file]
```

```
vparstatus -A [-M]
```

```
vparstatus -w
```

```
vparstatus -e
```

```
vparstatus -R [-p vp_name]
```



Example: vparstatus

- To show summary of all attributes and resources of all virtual partition:

```
# vparstatus
```

- To display detailed attributes and resources of partition vpar1:

```
# vparstatus -p vpar1 -v
```

- To display vpar2's resources in machine readable format:

```
# vparstatus -p vpar2 -M
```

- To see available resources that have not been assigned to any partitions:

```
# vparstatus -A
```



Boot Sequence

- With vPars, the monitor layer is added to the boot sequence
 - **ISL** **Initial System Loader**
 - **hpux** **Secondary System Loader**
 - **/stand/vpmon** **(vPars monitor and partition database)**
 - **/stand/vmunix** **(kernel of the virtual partitions)**



Example: Boot Sequence

- To load hpux with no partitions

ISL> hpux /stand/vmunix

- To boot vPars monitor

ISL> hpux /stand/vpmon <options>

NOTE: if no options are given, monitor will be in interactive mode

- To boot vPar from the monitor

MON> vparload -p <vparName>

- To boot additional vPars from a running vPar (vpar1)

vpar1# vparboot -p <vparName>

- *This can be executed from any running vPar*



Virtual Consoles

- **With limited number of PCI slots, you may not want to allocate one serial port for use as a console port for each partition you create.**
 - **With vPars, each virtual partition has its own virtual console.**
 - **First vPar created must own the Local Bus Adapter (LBA) that contains the physical hardware console port.**



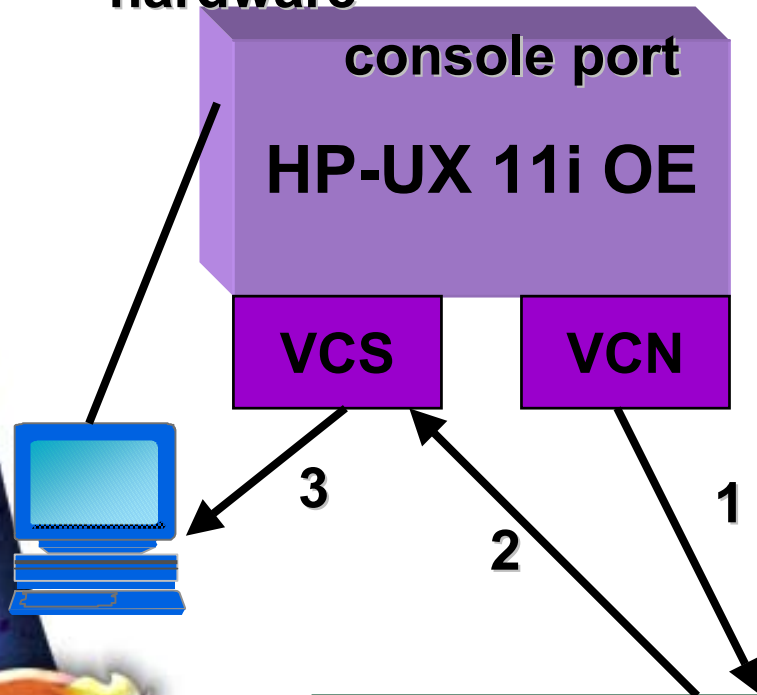
Virtual Consoles

- **Recommended way to access console**
 - **Default configuration: each vPar uses the virtual console**
 - *use Ctrl-A to cycle between all live vPars*
 - **Access h/w console through terminal, lan console, or lantronix**
 - **Alternative is to give one vPar the hardware console (no other vPar will have access to the console)**


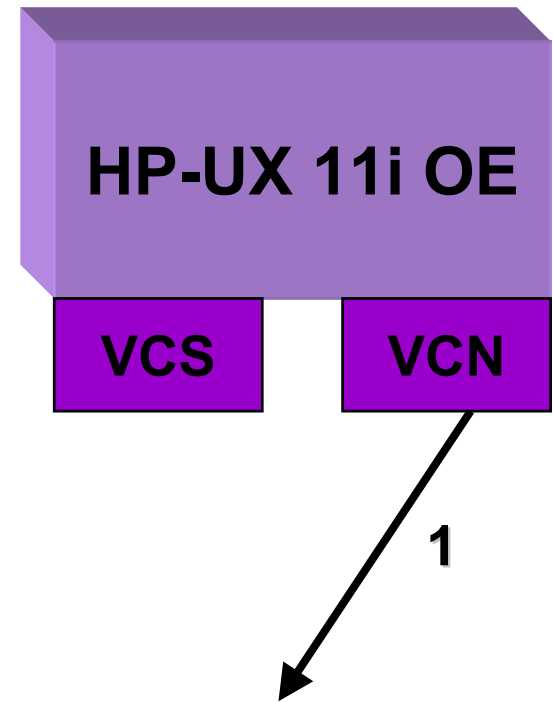


Virtual Console

**vPar 1- owns
hardware**



vPar 2



CPU Resources

Two types of CPUs:

- **Bound** : CPUs that are assigned to a partition when partition is active
 - *Interrupts are only handled by the bound CPUs*
- **Unbound**: CPUs that can be dynamically reassigned among partitions
 - **Provide greater flexibility of movement between virtual partitions**
 - Can be added and removed without needing to bring down the affected partitions.

Lowest numbered bound CPU is designated as the monarch CPU



When to use Bound or Unbound CPUs

Use bound CPU:

- When the application in the virtual partition are I/O intensive
- Only bound CPUs can process I/O interrupts
 - *With I/O intensive applications, there should be more bound CPUs than unbound CPU's*

Use unbound CPU:

- When the applications are CPU intensive
 - *Can easily adjust number of CPUs via dynamic CPU migration*



CPU Migration

- CPU Migration is adding and deleting CPUs from a virtual partition.
- Dynamic CPU Migration is migrating CPUs while the target virtual partition is running.
 - *Both bound and unbound CPUs can be migrated*
 - Only unbound CPUs can be dynamically migrated
 - HP-UX cannot dynamically migrate I/O interrupts. (bound CPUs)
- The upper limit on the number of CPUs that can be added when partition is running is the smaller of either:
 - *maximum count set for partition minus the bound count*
 - *available number of floaters*
- User has to balance the ratio of bound and unbound CPUs in a partition based on required I/O throughput



Memory Resources

- Three ways to specify memory:
 - Size in MB (recommended method)
 - *Starting address and size should be multiple of 64MB*
 - One or more ranges with each range containing starting byte address and size in MB (not recommended)
 - Combination of both
 - *A partition requires at-least 64MB below 2GB for kernel*
- Memory cannot be added or removed when partition is active



I/O Resources

- **System Bus Adapters (SBAs) and memory controllers are owned by monitor**
- **Local Bus Adapter (LBA) is the minimum granularity of assignment**
- **Combo SCSI/LAN card is not supported as a boot disk**
- **At the minimum a partition should own an LBA with boot disk and LAN port**
- **LBA cannot be removed or added to an active partition**
- **I/O path can have following attributes: console, boot, altboot, dump, and swap**
- **Full I/O path has to be specified when the path has attributes**
- **Monitor multiplexes console among partitions even if a partition owns the LBA that contains console**



Security Concerns

vPars commands restricted to root access

- Works on any partitions, regardless of which partition the commands are executed from.
- A user on one partition can affect another virtual partition by targeting the virtual partition in a vPars command.
 - For example, a root user running on the partition vpar2 can reset the partition vpar3 using the vparreset command.

With console access a user can gain access to the file systems on any of the virtual partitions in the same computer or hardware partition.

- This is done by performing a hardware reset of the computer (Ctrl-B), interrupting the boot sequence, and at the ISL prompt manually booting any partition into single-user mode.



Partition Plan

In the Partition Plan, the following has been identified:

- Number of virtual partitions to be created
- Virtual Partitions Names
 - *Each vPar requires a unique name*
 - *Maximum length of alphanumeric name is 239 characters*
- Hardware resource allocation
- Amount of memory to be allocated to each virtual partition
- Boot and LAN Paths has been determined
- Ensure hardware console port is owned by the first virtual partition



Partition Plan

vPar Name	vpar1	vpar2	vpar3
CPU	num = 3, min = 2	num = 2, min = 2	num = 2, min = 2
Memory	1GB	4GB	2GB
I/O Paths (LBAs)	0/0/* 0/4/*	0/8/* 1/10/*	0/5/* 1/4/*
Boot Path	0/0/2/0.6.0	0/8/0/0.5.0	1/4/0/0.5.0
Console	Virtual	Virtual	Virtual
Kernel Image	/stand/vmunix	/stand/vmunix	/stand/vmunix
Autoboot	On	On	On

Create a Partition Worksheet

Resource or Attribute	vparcreate Option
Partition name is vpar4	-p vpar4
Four total CPUs	-a cpu::4
Three bound CPUs and maximum of five CPU's.	-a cpu:::3:5
At Hardware Paths 101, 105, 107	-a cpu:101 -a cpu:105 -a cpu: 107
2048 MB Memory	-a mem::2048
All Hardware where path begins with 0/5	-a io:0/5
All Hardware where path begins with 1/4	-a io:1/4
Hardware at 1/4/0/0.5.0 as the boot disk	-a io:1/4/0/0.5.0:boot

vPars Commands

vPar commands:

- # **vparboot(1m)**
Boot a virtual partition
- # **vparcreate(1m)**
Create a vPar, with or without resources
- # **vparmodify(1m)**
Modify the attributes of a virtual partition
- # **vparremove(1m)**
Remove a virtual partition
- # **vparreset(1m)**
Reset a virtual partition
- # **vparstatus(1m)**
Display information about one or more vPars

vPar resources

- # **vparresources(5)**
- # **vpartition(5)**

vPars GUI

- # **vparmgr(1m)**



Example: vparcreate

- To create a virtual partition called vpar1 with no resources assigned:

```
# vparcreate -p vpar1
```

- To create a virtual partition called vpar2, setting the autoboot attribute and specifying five processors, two of which are bound, 2 Gb of physical memory, and a boot disk:

```
# vparcreate -p vpar2 -a cpu::5 -a cpu:::2:6 \  
-a mem::2048 -a io:1/0 -a io:1/0/2/0.6.0:BOOT\  
-B AUTO
```

- To set the primary boot path and the secondary boot path

```
# vparcreate -p vpar1 -a io:0/8/0/0.5.0:BOOT \  
-a io:0/8 -a io:0/8/0/0.2.0:ALTBOOT
```



Example: vparboot

- To boot the virtual partition vpar1:
`# vparboot -p vpar1`
- To boot the virtual partition vpar2 with an alternate kernel:
`# vparboot -p vpar2 -b /stand/vmunix_debug`
- To cold-install the virtual partition vpar3 from a file at the default path /opt/ignite/boot on the current virtual partition
`# vparboot -p vpar3 -I WINSTALL`



vparstatus Output

```
# vparstatus
```

```
[Virtual Partition]
```

Virtual Partition Name	State	Attributes	Kernel Path	Boot Opts
vpar1	Up	Stat,Auto	/stand/vmunix	
vpar2	Up	Dyn,Man1	/stand/vmunix	

```
[Virtual Partition Resource Summary]
```

Virtual Partition Name	CPU Min/Max	CPU Bound/ Unbound	Num IO devs	Memory (MB) # Ranges/ Total MB	Total MB
vpar1	2/ 4	2 1	3	2/640	704
vpar2	2/ 4	2 2	2	2/704	768



Installation Methods

The are several ways to install vPars:

- **As a cold install, along with an HP-UX 11i Server OE**
- **As an application on an HP-UX 11i Server OE**
- **From a running vPar, using the vparboot install process**
- **Using make_net_recovery (Ignite-UX)**



Initial vPars Install Process

- The process to install vPars on your HP-UX server is:
 - Choose a disk path to be the 'primary' boot disk for vPars
 - Install vPars software (T1335AC) on the vPars boot disk [use `swinstall(1m)`]
 - Verify vPars software installation
 - *Review swinstall log for errors or warnings*



Initial vPar Creation

To create the initial vPar

- Use the vparcreate(1M) command:

```
# vparcreate -p vpar1 -a cpu::2 -a cpu:::2 -a mem::640 \  
-a io:0/0 -a io:0/4 -a io:0/0/2/0.6.0:BOOT
```

- Do not assign the Monarch processor, if possible
 - *The Monarch is used by the monitor (/stand/vpmon)*
 - *The Monarch is usually the lowest numbered processor in the box or complex*
 - Use ioscan to verify



Additional vPar Creation

Creating additional vPars:

♦ Examples:

```
# vparcreate -p vpar2 -a cpu::1 -a cpu:::1:2 -a mem::1024 \  
-a io:1/2 -a io:1/4 -a io:1/2/0/0.6.0:BOOT \  
-a io:1/2/0/0.5.0:ALTBOOT
```

```
# vparcreate -p vpar3 -a cpu::1 -a cpu:::1:3 -a mem::8192 \  
-a io:X/Y -a io:X/Z -a io:X/Y/0/0.A.0:BOOT \  
-a io:X/Y/0/0.B.0:ALTBOOT
```

- ♦ It is not necessary to create additional vPars right now, but doing so may be more efficient.



vPar Installation Verification

- To verify and test vPar Installation:

- Shutdown HP-UX

```
# shutdown -ry 0
```

- Place the server in Boot Console Handler (BCH) mode

- Boot from primary disk, interact with ISL

```
Main Menu: Enter command or menu > bo pri Interact with IPL  
(Y, N, or Cancel)?> y  
ISL>
```



vPar Installation Verification

- To verify and test vPar Installation:
 - Boot the vPars Monitor “/stand/vpmon”

```
ISL> hpux /stand/vpmon
```

- From the Monitor, boot the first vPar

```
MON> vparload -p vpar1
```

- Access the vPar console

Login, and verify status

```
# vparstatus -v -p vpar1
```

•Copyright Hewlett-Packard Company 2003

Author: Asghar Riahi



vPars licensing & delivery

Base functionality

- Capability to create a maximum of 2 vpars per hard partition (or non-partitioned machine), where 1 vPar has only 1 CPU *{good for trial, upgrades, small machines}*
- Available at \$0 cost with HP-UX 11i (upon support for particular server)
- Available through HP Software Depot: www.software.hp.com

Full functionality (purchase T1335AC per CPU)

- Full functionality is licensed per installed CPU, in hard partitions (or non-partitioned machines) running vPars
- Available through separate vPars CD, upon order of product

Licensing impact on other HP OE or ISU software

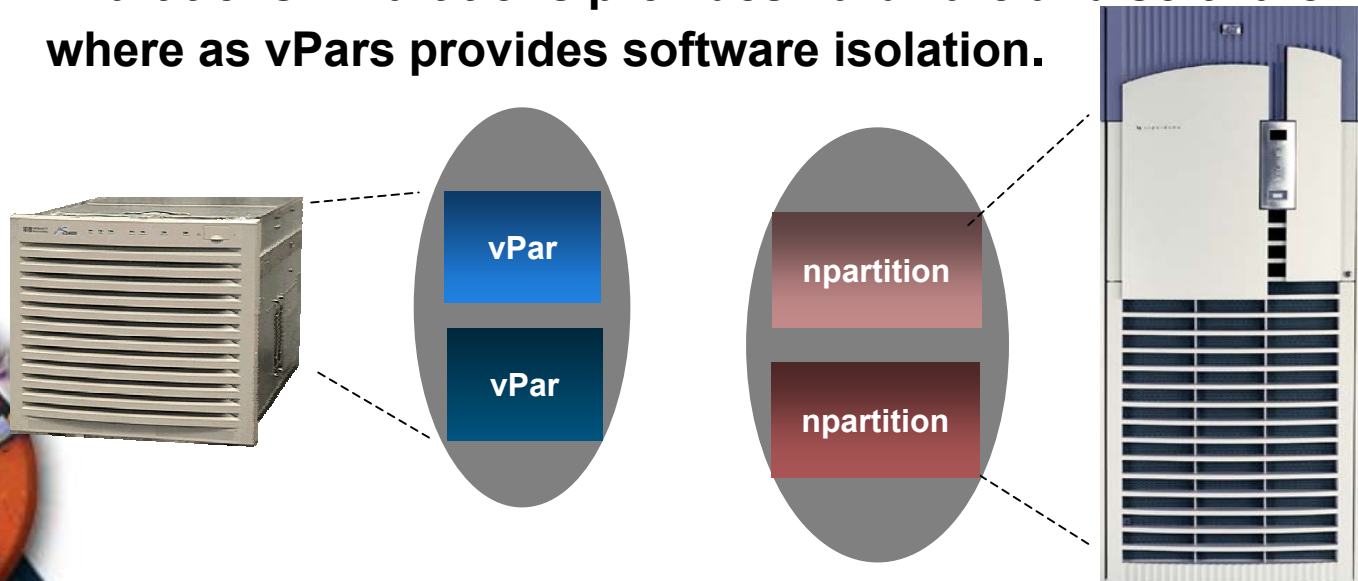
- None; HP OE or ISU software licensing per box, hard partition or tier is not impacted if that server becomes virtual-partitioned
OE and application software continues to be purchased on a hard partition basis *(since vPars can be so easily reconfigured)*



MC/ServiceGuard Cluster Configuration Considerations (a)


- ServiceGuard “cluster-in-a-box” is supported. Partitions of a cluster are wholly contained within the same hardware cabinet. While this configuration is subject to single points of failure, it may provide adequate availability characteristics for some applications.

Note: “cluster-in-a-box” consisting exclusively of vPars is more likely to suffer a complete cluster failure than one made up exclusively of nPartitions. nPartitions provides hardware and software isolation where as vPars provides software isolation.



MC/ServiceGuard Cluster Configuration Considerations (b)

- To configure a cluster with no SPOF, the membership must extend beyond a single hardware cabinet and comply to the standard ServiceGuard configuration rules. The cluster must be configured such that a failure of a single hardware cabinet does not result in the failure of a majority of the nodes in the cluster.
- Separate partitions within a physical server can be configured as part of different ServiceGuard clusters

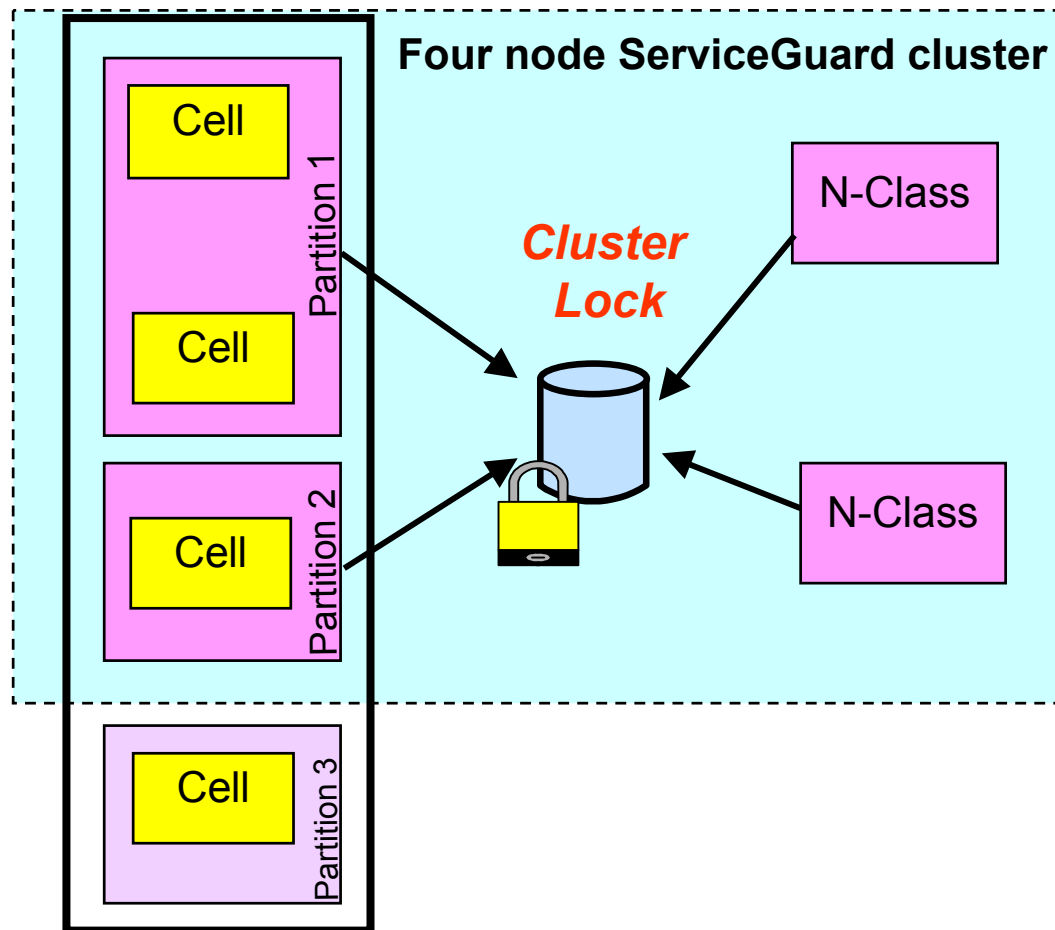


Mixed Clusters with 11.0 and 11i Supported

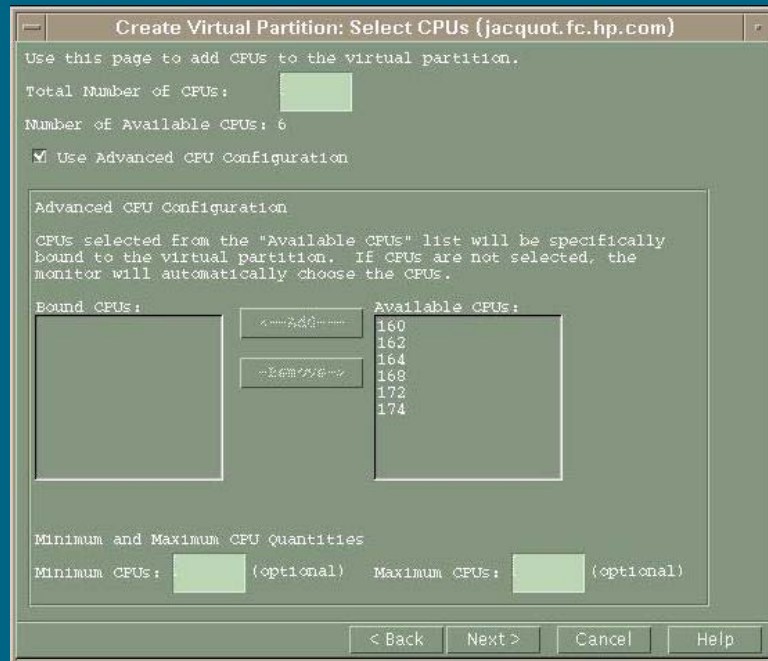
- Mixed clusters where some nodes are running HP-UX 11.0 and some nodes running HP-UX 11i is supported – requires all nodes running the same SG revision (e.g. 11.09 or 11.13)

- EXAMPLE where two nodes in a SuperDome are running HP-UX 11i and two N-class nodes are running 11.0

SuperDome System



virtual partition manager (vparmgr): GUI for managing virtual partitions



✓ *parmgr is vPar aware!*

(it doesn't do vPars configuration at this point, but the 2 are planned to be integrated in the future)

- Create, modify and delete virtual partitions (vpars)*
- Display assigned resources, attributes, and status of vpar
- Display vpar event log and samlog
- Boot and reset a vpar
- Direct invocation of task screens
- Preview create/modify vpar command lines prior to execution

Summary why use vPars?




To reduce total cost of ownership (TCO) for a server environment, due to:

- ✎ increased overall system utilization*
- ✎ increased flexibility in resource configuration*
- ✎ improved system availability (due to OS and application isolation)*
- ✎ improved scalability*

Production Environment

■ Partitioned ■ None Partitioned




Date: April 2003

Test and Development Environment

■ Partitioned ■ None Partitioned


None
Partitioned
25%



Partitioned
75%


Date: April 2003

KeyStone Partitioning



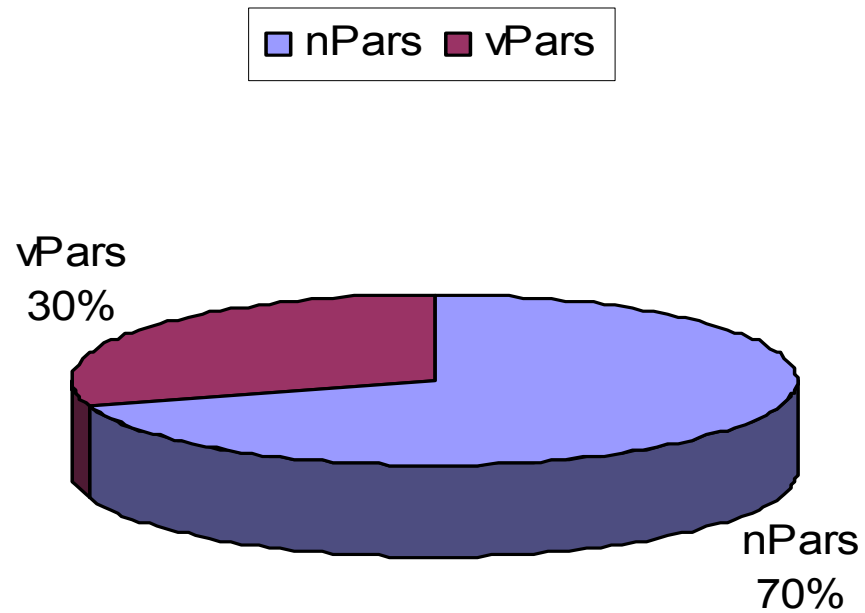
Date: April 2003

SuperDome Partitioning



Date: April 2003


Production environment



Date: April 2003

- External URLs

HP-UX Virtual Partitions



- www.hp.com/go/partitions
- then “HP-UX Virtual Partitions”
- www.hp.com/go/hp-ux
- www.software.hp.com
- www.docs.hp.com

