

# Introducing HP-UX 11i Virtual Partitions

A white paper from HP



Introduction .....	2
What are HP-UX 11i Virtual Partitions? .....	3
Why use HP-UX 11i Virtual Partitions? .....	4
Practical consolidation on enterprise-class servers .....	4
Increased utilization and scalability .....	5
More flexibility and capacity .....	5
Rapid deployment of new environments .....	6
Improved system availability .....	6
Improved performance and productivity .....	6
HP-UX 11i Virtual Partitions features .....	7
vPars resources .....	7
vPars and memory .....	7
vPars and I/O .....	7
vPars and console .....	7
Managing vPars .....	8
vPars commands .....	8
vPars with HP-UX 11i management tools .....	8
Compatibility with HP Virtual Server Environment and HP OpenView .....	9
VSE Partitioning .....	9
VSE Availability .....	9
VSE Utility Pricing .....	9
VSE Intelligent Control .....	9
HP OpenView Management Software .....	9
vPars operational overview .....	10
Virtual partition monitor (vPar monitor) .....	10
Availability .....	10
For more information .....	13

## Introduction

Today's data centers frequently need to be flexible in meeting the demands of new customers and new projects. Ideally, these demands are met quickly with minimal overhead while providing a relatively high level of isolation from other customers. For most data centers, the best solution is one that uses current capacity to satisfy these new demands, while continuing to provide existing customers with the capacity they expect. In doing so, data center owners can keep total cost of ownership to a minimum while maintaining a high level of responsiveness to their customers.

To address this challenge, Hewlett-Packard has created the HP Virtual Server Environment, an integrated server virtualization offering for HP Integrity and HP 9000 servers that provides a flexible computing environment maximizing usage of your server resources. The Virtual Server Environment encompasses a number of fully integrated, complementary components that enhance the functionality and flexibility of your server environment including workload management, availability software, partitioning and utility pricing.

HP offers the broadest range of partitioning capabilities as part of the HP Partitioning Continuum. These capabilities are offered through a combination of hard partitions, soft partitions, and resource partitions, which can be used together for even greater combined functionality.

This paper introduces HP-UX 11i Virtual Partitions (vPars) – soft partitions that are a part of the HP Partitioning Continuum for HP-UX. Virtual partitions can provide granularity to the hardware level (CPU core, I/O card) with each vPar running its own version of the HP-UX 11i operating system. This is a good approach when applications need fewer resources than in an individual server or nPartition, and the applications need to be isolated from one another for reliability, tuning, or version compatibility of the software stack.

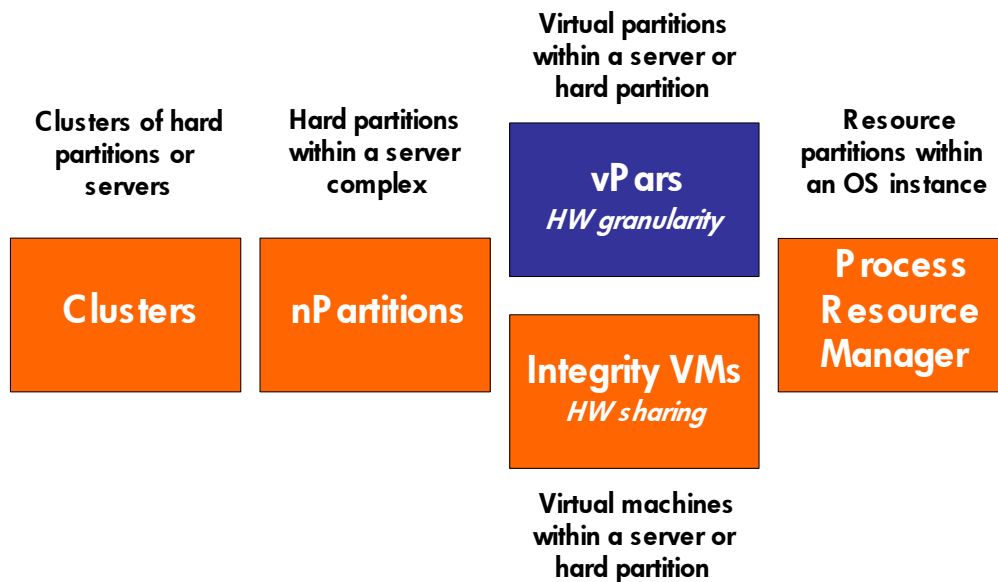


Figure 1 HP Partitioning Continuum for HP-UX

## What are HP-UX 11i Virtual Partitions?

For effective partitioning, it is necessary to isolate operating environments such that multiple customers' applications can coexist in the same server or cluster, while assuring complete privacy. In addition, it is often necessary to dynamically create, modify or even delete the isolated operating environments on a running server *without interrupting non-related partitions*. To meet this need, HP has developed HP-UX 11i Virtual Partitions (vPars) - a unique technology that provides application (including name space) and operating systems isolation.

The virtual partitions can be dynamically created using software commands. Each partition runs its own image of the HP-UX 11i operating system and can host its own applications in a fully software isolated environment. Within each virtual partition, up to 64 resource partitions can be created and utilized using solutions such as Process Resource Manager (PRM) and HP-UX Workload Manager (WLM).

Figure 2 is an example of a single rp7410 server running four vPars.




Figure 2 – Virtual Partitions within a Single rp7410 Server

## Why use HP-UX 11i Virtual Partitions?

Most data centers have server hardware that is underutilized. These same data centers almost never have enough servers to satisfy their customers. These seemingly contradictory situations have led administrators to the obvious conclusion for years – “I should be able to make better use of the resources I already have.”

HP-UX 11i Virtual Partitions is a powerful, flexible tool that makes it possible to run multiple workloads – each with their unique OS configuration requirements – on the same server at the same time. With vPars, customers can carve their hard partitions (nPartitions) or servers into multiple virtual servers, applying unutilized resources to additional virtual partitions. CPUs can be migrated between those virtual servers, as needed (through commands, scripts or WLM.)

Often customers combine the the availability of HP’s hard partitions (nPartitions), with its electrical isolation, with the increased flexibility and granularity of resources provided by HP soft partitions (HP-UX 11i Virtual Partitions) to make an even more powerful partitioning solution.

vPars enables:

- Practical consolidation on enterprise-class servers
- Increased utilization and scalability
- More flexibility and capacity
- Rapid deployment of new environments
- Improved system availability
- Improved performance and productivity

## Practical consolidation on enterprise-class servers

There are pros and cons of consolidation in the data center. Approaches to server consolidation are varied and can be achieved at many levels:

- Application
- Middleware
- Operating System
- Hardware

Consolidation at any of these levels implies some form of consolidation at lower levels. An example of consolidating at the middleware level is reducing the number of Oracle database instances. It is difficult to combine these instances unless they reside on the same OS instance. The Oracle instances are no longer running on separate OS instances nor do they reside on separate physical servers. Therefore, hardware and OS consolidation has already taken place either explicitly or implicitly.

With consolidation, as with almost any endeavor, 20 percent of the effort gives 80 percent of the benefit. In comparison to other consolidation approaches, consolidating at the hardware level is analogous to 20 percent of the effort. Hardware consolidation with HP-UX 11i Virtual Partitions enables you to realize most of the consolidation benefit.

Consolidating at other levels requires additional effort for diminishing returns. Application consolidation can be time-consuming and expensive. For example, moving all users to a single mail application, can take months, depending on the organization. It can also have a high level of impact

on productivity and morale if not done with much careful planning and support. Similarly, consolidating middleware instances requires a high level of expertise in the middleware, an understanding of how it is being used by the various business applications, and what performance impacts may result from the consolidation. Mutually exclusive OS or middleware version requirements (or both) frequently block consolidation of OS instances. Within a single OS instance, software may be incompatible. In fact, some Independent Software Vendors (ISVs) only provide support on dedicated servers.

Hardware consolidation, that is, multiple “virtual” servers on a single physical server, avoids many of the pitfalls of consolidation at higher levels.

vPars provides a great hardware consolidation solution. It enables multiple virtual partitions, with unique OS and application requirements and combinations, to share the same physical server.

vPars provides a unique combination of a practical consolidation solution with high-performance, enterprise-class hardware.

## Increased utilization and scalability

Very few data centers achieve utilization rates above 40 percent. This is usually the result of two common practices: First, there is typically a one-to-one relationship between applications and servers, which rarely fully utilizes the system. Second, many enterprise-class applications are resource-intensive, but the high demand is only for short periods. Many businesses have applications that are cyclic in nature, such as payroll and end-of-month billing. Frequently, the cycles for the applications do not coincide. Servers are normally configured to handle the peak loads, causing lower average utilization of the server.

Virtual partitions can be used to improve the ratio of applications to physical resources since:


- 1) additional vPars can be created to utilize unused resources
- 2) CPU resources can be moved between vPars when and where required, without having to reboot the affected partitions

Each virtual partition is allocated: a specified amount of memory, one or more CPUs, and one or more I/O cards. Hence, multiple, isolated virtual servers, with separate instances of the HP-UX 11i operating system, applications and users, run on the same physical server. This is practical for enterprise-class and mission-critical applications that require the scalability inherent in symmetric multi-processor servers (SMPs).

HP-UX 11i Virtual Partitions can dynamically partition hardware, and dedicate resources to those virtual partitions. This provides flexibility in scalability of the individual partitions, since one or more CPUs can be allocated to each partition. Figure 3 shows this approach with hardware dedicated to each partition.

---

*vPars increases system utilization, while providing good partitioning scalability*



**Figure 3.** vPars – system utilization and scalability

---

## More flexibility and capacity

HP-UX 11i vPars enables greater flexibility in configuring servers. For example, many businesses have applications that are cyclic in nature such as, payroll, end of month billing etc. Many times, the cycles for the applications are not aligned. Normally, servers need to be configured to handle peak loads. This can lead to poor utilization. With vPars, applications can be allocated to a larger percentage of the system resources during the peak usage times, and utilize less resources during off-peak times. This frees up hardware resources that can be assigned to other vPars where the applications may be experiencing a high demand for these resources.

In addition, one of the inherent problems in a single system is the difficulty in expanding CPU resources when the demands of the application, or multiple applications, exceed the server's configuration. Usually the system would need to be shut down and additional CPUs added. vPars CPU cores can, however, be dynamically moved from one vPar to another, without bringing the system down, through either vPars commands, or automatically, through the use of HP-UX Workload Manager. This allows CPU resources to be moved to vPars with the greatest workloads, or removed from vPars when they are no longer required.

## Rapid deployment of new environments

vPars' flexibility is complemented by its ability to rapidly deploy a new virtual server. Instead of purchasing new servers for new projects, unused server resources can be utilized to create additional vPars housing new environments. This eliminates expensive and time-consuming server acquisitions for new deployments.

Users with privileged access can create, deploy, monitor, and perform overall virtual partition management from any vPar.

In addition, virtual partitions enable SMP software development and testing for multiple configurations of virtual SMP servers – all on one host server. On the other hand, prototyping and capacity testing of distributed applications is accommodated by using dozens of HP-UX 11i Virtual Partitions on a single nPartition or physical server. Replication of the OE and application environments from one to another vPar is made easy through IgniteUX.

## Improved system availability

In traditional server environments, all CPUs within a server run the same OS instance and one or more applications. Application and OS failures may affect the entire server.

With HP-UX 11i Virtual Partitions, however, the vPar monitor sits logically between the hardware and each OS, allowing each partition to support a different instance of HP-UX 11i. HP-UX 11i Virtual Partitions house completely separate operating system instances isolated by software. Each vPar is given its own CPU, memory and I/O resources.

For improved single system availability, you can run one application per Virtual Partition. When an application software failure does occur in one of the HP-UX 11i Virtual Partitions, that application may be lost, but the applications on other HP-UX 11i Virtual Partitions continue to run. In fact, even if the OS panics in one of the vPars, applications running on the other vPars are not affected. In this way, running HP-UX 11i Virtual Partitions limits the impact of application or OS failures on overall OS and application availability.

Serviceguard can be used to further improve application availability.

## Improved performance and productivity

Since HP-UX 11i Virtual Partitions allocates resources to the different vPars, those resources have isolated use for a specific vPar. In addition, specific OS instances can be tuned for a specific application's performance.

Sometimes, customers - for example, ISPs, deploy small Web servers for customers that need to manage their own content. The importance of the deployment is to maintain customer data isolation in addition to performance or application-specific tuning. vPars can provide a means for isolating one set of users from another.

vPars can also be used to set up isolated partitions as test environments. This could be for new revisions of current applications, or for deployment of new applications. vPars allows testing on the exact deployment environment. This improves the quality of the test without incurring the costs of replicating the deployment environment.

## HP-UX 11i Virtual Partitions features

- Support for many virtual partitions per nPartition or server (each with their own instance of HP-UX 11i and their own applications)
- Runs on single, uni-processor and SMP, select HP9000 or HP Integrity servers, or within HP nPartitions on such servers
- Hardware granularity (single CPU core, single I/O card)
- Dynamic reallocation of physical CPU cores across virtual partitions
- Software fault isolation (application – including name space – and OS isolation)
- Individual reconfiguration and reboot
- Single Shared (Multiplexed) System Console
- Compatible with the other HP VSE solutions and availability software (e.g. HP Process Resource Manager (PRM), HP-UX Workload Manager (WLM), HP Integrity Essentials Global Workload Manager (gWLM), Systems Insight Manager (SIM), and Serviceguard, etc.)
- Fully supported and serviced by HP

## vPars resources

### vPars and memory

Allocation of memory is specified in blocks of 64MB (vPars A.03+) and "granularity" (vPars A.04+). Sharing of physical memory is not allowed between multiple vPars. Each vPar requires a minimal amount of memory for booting HP-UX 11i and running the applications.

### vPars and I/O

vPars requires that each Local PCI Bus Adapter (top-level PCI bus) is assigned to, at most, one vPar. Therefore, the Local Bus Adapters (LBAs), and the interface cards attached to those LBAs, may not be shared among vPars. This includes multi-port and multi-function cards. They cannot be split between vPars.

## vPars and console

The system console on a standalone (non-partitioned) server serves two distinct and unique purposes. First, the system console is used to monitor and interact with the Guardian Service Processor (GSP) and Initial System Loader (ISL). Additionally, the system console is used to monitor and interact with HP-UX, particularly when the system is in single user mode or when networking or terminal services are unavailable.

With vPars, the system console continues to be used for interacting with GSP and ISL, and the vPar to which it is physically attached. Other vPars use a “virtual console.” Together, the vPars use the system console in a multiplexed fashion. When the system console is multiplexed among several vPars, a system administrator with access to the system console may open a console session on any of those vPars using a keystroke (Ctrl-A) at the console.

## Managing vPars

The definitions of the vPars are stored as a binary file on all of the vPar boot disks. The default location is /stand/vpdb on the disk that the monitor image was loaded from – usually the system boot device. Changes to the vPars configuration are made using a special set of commands. These commands modify the database in memory, which is synchronized with all of the vPar databases on disk and preserved across reboots. The commands can be issued from any of the active vPars, provided appropriate user permission is granted.

### vPar commands

vPar commands are available to the system administrator. These commands can be executed from any active vPar, or in some cases, to create the initial vPar/database. A summary is as follows:

**vparcreate** – create a vPar, with or without resources

**vparremove** – destroy a vPar, where the partition definition is removed and all resources associated with the given vPar are reclaimed by monitor

**vparmodify** – add or modifies attributes of an existing vPar (such as changing the boot device)

**vparboot** – start/boot a vPar

**vparreset** – stop/reset a vPar

**vparstatus** – display information about one or more vPars

## vPars with HP-UX 11i management tools

vPar management can be performed via command line interface. With the few vPar commands, and the ease of using them in UNIX scripts, vPars can be integrated into the many industry recognized system management tools to help “automate” the administration of vPars.

The vPar administrator can perform the following tasks:

- Create a new virtual partition.
- Modify an existing virtual partition.
- Delete an existing virtual partition.
- Display the resources assigned to each virtual partition.
- Display the status of each virtual partition.
- Display logs of vPars activity.



- Boot a virtual partition.
- Reset a virtual partition.
- Protection against applications in one vPar from intentionally or accidentally reading or writing another vPar's memory (a vPar only aware of its own memory)
- Resources (such as CPU or Memory) cannot be accessed simultaneously by two vPars. Resources must first be removed from one vPar then added to another
- Non-root users cannot create, modify or destroy vPars
- A hardware resource must be made available before it can be added to a vPar
- vPar commands can run simultaneously from multiple vPars, but all access to the monitor will be serialized
- The monitor ensures that the on-disk copies of the partition database are synchronized with the monitor's internal state

Each virtual partition functions like a standalone server and thus is administered by the root user like any other HP-UX server. vPar commands to modify vPars are performed by root. The root user on any vPar has the flexibility to run vPar commands for its vPar or any other vPar on the system.

## Compatibility with HP Virtual Server Environment and HP OpenView

HP-UX 11i Virtual Partitions (vPars) are an important component of the HP Virtual Server Environment for HP-UX 11i. vPars are well integrated with many of the other elements of the Virtual Server Environment as well as with many of HP's extensive line of OpenView enterprise management products, a few of which are highlighted here.

### VSE Partitioning

vPars is compatible with other partitioning elements within the Virtual Server Environment. vPars are supported running on nPartitions. Intra-OS resource partitioning for HP-UX 11i, such as that offered by PRM groups, both processor sets and sub-CPU groups, also work within vPars. While vPars and HP Integrity Virtual Machines are not run simultaneously in the same hard partition, they can run on the same server in separate hard partitions.

### VSE Availability

HP Serviceguard is supported with HP-UX 11i Virtual Partitions, providing high availability of applications on virtual partitions. With Serviceguard, an application that fails in a virtual partition on one host will automatically start in another virtual partition on the same host, a virtual partition hosted by another host, or another physical server not running HP-UX 11i Virtual Partitions.

Special care needs to be taken when designing the highest availability within a vPar environment. Because of the various vPar configurations on HP 9000 or Integrity servers, with or without hard partitions, careful design criteria need to be followed. The Serviceguard Support Organization has supplied recommended guidelines when using vPars in mission critical implementations.

### VSE Utility Pricing

Utility Pricing is an option available on many of the HP9000 and HP Integrity servers running HP-UX 11i. With Instant Capacity or Temporary Instant Capacity, you can purchase a fully CPU populated server up front, but only activate a sub-set of the total CPUs. Instant Capacity allows you to purchase

the remainder of the CPUs as you require them. Temporary Instant Capacity allows you to purchase CPU time to activate and deactivate CPUs as required by the workloads.

Another alternative is Pay-per-use. This allows customers to pay for their server resources partially based on a right to access fee, and partially based on actual usage.

vPars works seamlessly with all utility pricing options. Additional CPU capacity can be turned on or off and added or removed from vPars dynamically, without a reboot.

## VSE Intelligent Control

While the partitioning, availability, and utility pricing components of the HP Virtual Server Environment provide a flexible environment, that is just the beginning of the story. The VSE goes beyond creating just a flexible environment by automating it. The intelligent control products are designed to work with the flexible environment, including vPars, and help automate it.

HP Integrity Essentials Global Workload Manager and HP-UX Workload Manager are the automation tools in the VSE. They contain the business policies for the environment, measure current performance or utilization, and automatically move CPU resources among vPars or other types of partitions based on the actual resource demands. These products also automate the activation and deactivation of utility pricing components.

The HP Integrity Essentials products support not only HP-UX 11i on HP Integrity Servers, but also HP-UX 11i on HP 9000 servers.

The HP Virtual Server Environment for HP-UX 11i, of which vPars is a key part, creates a flexible environment that can double server resource utilization, maintain continuous service levels, and provide a mechanism to pay only for the resources that are used. In short, it creates a more agile and lower cost environment.

## HP OpenView Management Software

HP OpenView tools such as GlancePlus can monitor or manage individual HP-UX 11i Virtual Partitions. Using GlancePlus on the nPartition or server reflects the physical resource utilization for all HP-UX 11i Virtual Partitions running on an nPartition or server. When run in a virtual partition, GlancePlus monitors resource utilization of the CPUs and I/O for the operating system running in that VPAR.

## vPars operational overview

To understand how vPars works at a high level, it is best to compare it to a generic HP-UX server. Figure 8 shows a 4-way HP-UX server. Without vPars, the entire server would run a single instance of HP-UX. All of the resources (CPU, Memory and Disk) would be dedicated to the applications running

in this single instance. The software stack for this server would look like the one in Figure 9.

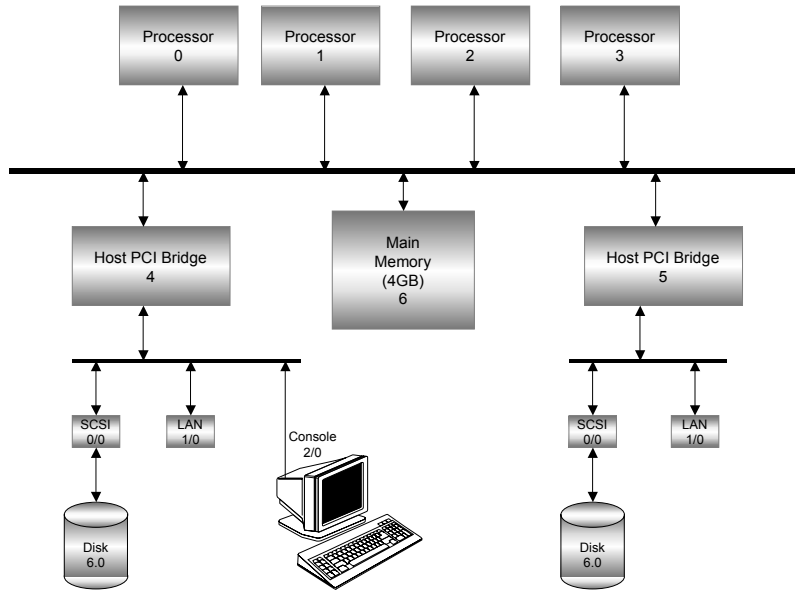


Figure 4 – Generic HP-UX server block diagram




Figure 5 – Generic HP-UX server software stack

Using vPars, the server in Figure 9 can be broken into two partitions, each with a subset of the hardware (Figure 10). Each vPar has its own boot disk, at least one CPU and one LAN connection, and enough memory to run HP-UX and the applications that will be hosted on this vPar. Since each vPar can run its own copy of HP-UX (even at different patch levels), each is completely isolated from software errors, system panics, etc. A software stack for two vPars is shown in Figure 11.

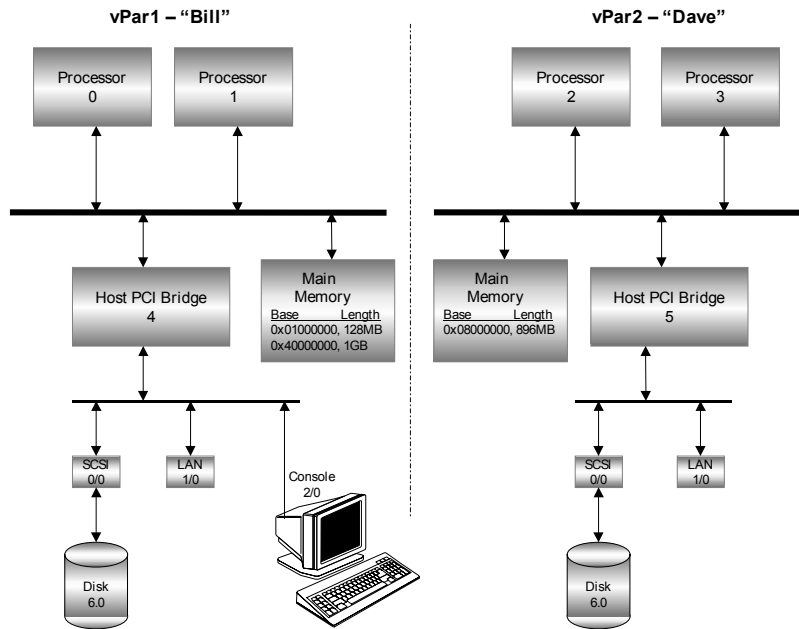


Figure 6 – Generic HP-UX server block diagram with two vPars




Figure 7 – Software stack with two vPars

## Virtual Partition Monitor (vPar monitor)

Notice that there is an additional layer of software and/or firmware in the vPar solution stack, the Virtual Partition Monitor (vPar monitor).

The monitor manages the partitioning of the resources, loads kernels, and emulates global platform resources to create the illusion that each instance of HP-UX that it is on a standalone system with only the resources that have been dedicated to that vPar.

Each instance of HP-UX is completely unaware of the other hardware in the system. The individual instances of HP-UX have complete ownership of the hardware resources that are assigned to them.

The monitor is not involved in accessing I/O hardware or physical memory once it has transferred ownership of the hardware to a vPar.

At the heart of the monitor is the partition database, which tracks the resources that are associated with each vPar. When the vPar monitor is running, the master copy of this database is kept by the monitor in memory. Changes to vPars through the use of vPar commands are first verified by the monitor, then applied to the database in memory. This in-memory database is then synchronized out to the copies on disk. Because the database is stored on disk, all changes to the partition database are preserved across system reboots.

In a system running without vPars, the HP-UX kernel (vParunix) is booted by the secondary loader (hpux) from the Initial System Loader (ISL) prompt using a command like:

```
ISL> hpux /stand/vParunix
```

If vPars are used, the vPar monitor (vpmon) is booted from ISL instead of the HP-UX kernel. The monitor prompt (MON>) would then be available from which to load the individual vPars. In this case, this initial load command would look like:

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

```
MON>
```

The monitor code is loaded from the file `/stand/vpmon` on the system boot device in the same way as a normal HP-UX kernel would be loaded from the file `/stand/vParunix`. The monitor loads the partition database from `/stand/vpdb` and creates vPars according to the resources allocations in the database. If there are no command-line options, the monitor is booted in interactive mode with a command-line interface. Once the vPars are up and running, the vPars Monitor is infrequently invoked. In fact, the monitor is invoked only when HP-UX makes calls to firmware, when the OS is shutting down, or when vPars management commands are executed.

## Availability

The HP-UX Virtual Partitions product is available on cell-based HP9000 servers (Superdomes, rp84xx, rp7410, rp7405), with or without nPartitions, as well as on the single HP9000 rp7400/N4000, rp5470/L3000 and rp5405 servers, running HP-UX 11i v1 or HP-UX 11i v2. vPars also supports cell-based Integrity servers (Superdome, rx8620, rx7620), with or without nPartitions, running HP-UX 11i v2.

## For more information

For additional information, please visit the following sites:

- HP Partitioning Continuum  
[www.hp.com/go/partitions](http://www.hp.com/go/partitions)
- HP Virtual Server Environment  
[www.hp.com/go/vse](http://www.hp.com/go/vse)
- HP-UX 11i  
[www.hp.com/go/HP-UX11i](http://www.hp.com/go/HP-UX11i)
- HP Integrity Servers  
[www.hp.com/go/integrity](http://www.hp.com/go/integrity)

- HP 9000 Servers  
[www.hp.com/go/hp9000](http://www.hp.com/go/hp9000)
- vPars System Administrator's Guide: Installing and Managing HP-UX Virtual Partitions  
<http://docs.hp.com/hpux/11iv1> , or <http://docs.hp.com/hpux/11iv2>, under "Virtual Partitions"

© 2005 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the U.S. and other countries. Oracle is a registered U.S. trademark of Oracle Corporation. UNIX is a registered trademark of The Open Group.

26August 2005

