



# Realize new workload migration and consolidation possibilities

Upgrade your vPars to HP-UX vPars and Integrity VM v6

## Table of contents

<b>Introduction</b>	<b>3</b>
Intended audience	3
<b>Glossary</b>	<b>3</b>
<b>An overview of HP-UX vPars and Integrity VM 6.1</b>	<b>4</b>
About HP-UX Virtual Partitions	4
About HP Integrity Virtual Machines (VMs)	4
New features and enhancements	4
<b>Assumptions</b>	<b>5</b>
<b>Prerequisites and limitations</b>	<b>5</b>
Software dependencies	5
Hardware dependencies	5
<b>Process overview</b>	<b>6</b>
<b>Prepare for the migration</b>	<b>6</b>
Review the new features and paradigms	6
Review the Virtual Partition command changes	7
Plan the new system	8
Extract data from the existing configuration	8
Review the configuration and verify the new system’s hardware requirements	9
<b>Install the VSP (server or nPartition) with HP-UX vPars and Integrity VM 6.1</b>	<b>9</b>
<b>Generate commands for creating the new vPars and infrastructure</b>	<b>10</b>
Examples	10
Verifying the new vPar and VSP configurations	11
<b>Provision virtual partitions</b>	<b>14</b>
Starting with fresh boot images	14
Migrating disks from legacy servers	14

<b>Post migration steps</b>	<b>16</b>
Confirm successful vPar boot and data visibility	16
Update applications and I/O paths as necessary	17
<b>Conclusion</b>	<b>18</b>

# Introduction

The Mission-Critical Converged Infrastructure is the foundation for the next decade of computing. The newest line of HP Integrity systems combines years of trusted HP Integrity resiliency with HP BladeSystem efficiencies. The HP mission-critical Converged Infrastructure creates an ecosystem with four key objectives:

- **Simplify and unify IT:** with a common, modular architecture from x86 to Superdome 2
- **Always-on resiliency:** a secure and reliable infrastructure from CPU to solution
- **Dynamic optimization:** integrated management and virtualization to optimally scale resources
- **Investment protection and stability:** sustained innovation, decades of support life, and compelling value

Specifically, this document is intended to provide processes and pointers for those administrators who are currently running HP-UX Virtual Partitions on legacy, cell-based systems (rp74nn, rp84nn, rx76nn, rx86nn, Superdome) and who are considering upgrading to HP-UX vPars version 6.1. A brief set of feature descriptions, benefits and limitations, and configuration recommendations are provided, as are step-by-step instructions on how to perform the migration.

It is recommended that the **HP-UX vPars and Integrity VM 6.1 administrator guide and release notes** be read before performing the migration.

## Intended audience

This document is intended for system administrators responsible for setting up and managing HP-UX vPars. Administrators should be familiar with the product as well as how to install and manage HP-UX operating environments.

## Glossary

AVIO: Accelerated Virtual I/O

CLI: Command Line Interface

DC-OE: HP-UX 11i v3 Data Center Operating Environment

DIO: Direct I/O

DRD: Dynamic Root Disk

EFI: Extensible Firmware Interface

GUI: Graphical User Interface

GUIDMGR: GUI ID manager

HBA: Host Bus Adapter (I/O card)

NPIV: N-Port ID Virtualization

OE: Operating Environment

vPars: Virtual Partitions

VSE: Virtual Server Environment

VSE-OE: HP-UX 11i v3 Virtual Server Environment Operating Environment

VSP: Virtualization Services Platform

# An overview of HP-UX vPars and Integrity VM 6.1

HP-UX vPars and Integrity VM 6.1 is a combined product that enables you to run either Virtual Partitions version 6.1 (vPars v6.1) or Integrity VM version 6.1 (but not both at the same time). For the purpose of this document, it is assumed that the migration being considered is from “legacy” vPars, either the firmware-based Superdome 2 virtual partitions (v5), or software-based v5 vPars on cell based systems to vPars v6.1.

## About HP-UX Virtual Partitions

vPars v6.1 is an HP-UX Virtual Partitions product that enables you to run multiple instances of HP-UX simultaneously on a Virtualization Services Platform (VSP). Each virtual partition is assigned a subset of hardware, runs a separate instance of HP-UX, and hosts its own set of applications. The VSP virtualizes physical processors, memory, and I/O devices, allowing you to allocate them as virtual resources to each virtual machine. As each instance of HP-UX is isolated from all other instances, vPars provides application and Operating System (OS) fault isolation. Applications running on top of HP-UX (using vPars), run the same as when run on HP-UX native mode (standalone). No changes, recompilation, or re-certification is required for applications running in the virtual partitions (unless otherwise noted). The patch level of each instance of HP-UX can be different. vPars v6.1 offers a common manageability framework that is accessible from the VSP. You can use the VSP to assign resources that includes CPU, memory, and I/O, and manage virtual partitions from the Integrity Virtual Server Manager graphical user interface (GUI). The VSP provides control for creating, starting, stopping, modifying, and deleting virtual partitions.

## About HP Integrity Virtual Machines (VMs)

Integrity VM is a virtualization technology that provides operating system isolation, with sub-processor core allocation and shared I/O. Integrity VM can be installed on an Integrity server, Integrity server blade, or hardware partition (nPartition) running HP-UX. The environment consists of two types of components:

- VSP
- VMs

Virtual machines are abstractions of real, physical machines. The guest OS runs on the VM just as it would run on a physical Integrity server, with no special modification.

Guests are fully-loaded operational systems, complete with OS, applications, system management utilities, and networks, all running in the VM environment. You boot and manage guests using the same storage media and procedures that you would if the guest OS were running on its own dedicated physical hardware platform. Even the system administration privileges can be allocated to specific VM administrators.

Because multiple VMs share the same physical resources, I/O devices can be allocated to multiple guests, maximizing use of the I/O devices and reducing the maintenance costs of the data center.

## New features and enhancements

HP-UX vPars v6 adds C-Class Integrity blade support, as well as a fully functional GUI that can create, modify and delete vPars as well as monitor, start, and stop them. The management framework, including the GUI, is leveraged from Integrity VM and provides greater levels of I/O sharing and cost savings.

The new I/O paradigm enables you to choose I/O on a port-by-port basis, and how much to share the I/O hardware between vPars (if at all). Where speed and bandwidth are key, the port can be dedicated to a single vPar and can use technologies such as N-Port ID Virtualization (NPIV) and Direct I/O (DIO). Where I/O performance is not the key consideration, it is possible to share each I/O port. This allows the server to increase the number of vPars without purchasing new adapters or being limited by the number of I/O slots on the physical server.

## Assumptions

- This is a legacy vPars to HP-UX vPars v6.1 upgrade; no VMs are involved
- The upgrade is to optimize for low overhead, occasionally at the expense of flexibility
- These assumptions imply that Fibre Channel ports will be configured using NPIV as virtual HBAs (vHBAs) wherever possible.

## Prerequisites and limitations

### Software dependencies

The software dependencies for HP-UX vPars and Integrity VM 6.1 are split between the VSP, which services and manages the virtual environments, and the virtual partitions themselves.

#### VSP dependencies

The VSP software is supported on the HP-UX 11i v3, March 2012 and later update releases. It is included with the HP-UX 11i v3 VSE OE and Data Center OE and may also be purchased separately. The bundles are:

- BB068AA      HP-UX vPars & Integrity VM v6
- GUIDMGR      HP-UX GUID Manager (optional, recommended if NPIV is desired)
- VirtualBase      Base virtualization software

These bundles do not include the Integrity Virtual Server Manager (VSM GUI), an HP System Management Homepage (SMH) plug-in. The GUI is a separately installed bundle (vmmgr) and is available as part of the OEs. It may also be downloaded from the HP Software Depot at [software.hp.com](http://software.hp.com).

#### Virtual partition dependencies

The virtual partitions must run the HP-UX 11i v3, September 2011 update release or later. Any OE is acceptable; however, the vPar's OS must also have:

- VirtualBase      Base Virtualization Software
- PHKL\_42464 or a superseding kernel patch.

### Hardware dependencies

All Integrity Superdome 2 nPartitions, BL8x0 i2 blades, and rx2800 i2 servers are supported. The nPartitions where the vPars will reside must have sufficient cores, memory, and I/O to support both the virtual partitions and the VSP.

If NPIV is desired for storage (recommended), then Fibre Channel cards and infrastructure that support this virtualization will be required. See the latest supported hardware lists at the HP-UX Virtual Partitions manual website at [hp.com/go/hpux-vpars-docs](http://hp.com/go/hpux-vpars-docs) or the latest release notes for an up-to-date list.

If DIO is desired for networking (recommended), then network cards that support DIO will be required. See the latest supported hardware lists at [hp.com/go/hpux-vpars-docs](http://hp.com/go/hpux-vpars-docs) or the latest HP-UX vPars and Integrity VM 6.1 release notes for an up-to-date list.

## Process overview

- Review the new capabilities available on HP-UX vPars v6.1 and their associated resource requirements
- Design the new VSP (physical server) to meet your current and anticipated needs
- Install, provision, and configure the VSP
- Generate and run vPar creation scripts (or create the vPars with the Integrity Virtual Server Manager GUI)
- Install or migrate the virtual partition boot disks
- Install or update applications
- Mount/copy existing data disks
- Validate and tune the environment

## Prepare for the migration

### Review the new features and paradigms

It is strongly recommended that you review the section titled “**An overview of HP-UX vPars and Integrity VM 6.1**” in this document and the HP-UX vPars and Integrity VM 6.1 Release notes. There are several new features that you may wish to take advantage of which will impact your server configuration. For instance, NPIV is only supported by some of the available Fibre Channel cards. Likewise, DIO for networking is supported only by a subset of the available network interfaces (NICs).

You should also determine whether you are going to try to preserve your existing boot disk. This is not generally recommended, but may be possible for SAN-based boot devices that have been updated within the last 2–3 years. It is not possible to reuse boot devices from PA-RISC systems.

Once you have decided which capabilities are desirable, refer to the following table for the resulting configuration impacts.

**Table 1.** Capability requirements

Feature	Requirements
<b>NPIV</b>	NPIV capable HBAs include 451871-B21, 403619-B21, AH400A, AH401A, AB378B, AB379B, AD193A, AD194A, and AD300A. More HBAs may have been added since the writing of this paper.
<b>DIO</b>	DIO capable NICs include: <ul style="list-style-type: none"><li>• SD2 &amp; BL8x0 i2 LOMs</li><li>• NC532m, NC552m, and NC553m mezzanine cards</li><li>• AM225A/232A/233A, AT111A, and AT118A PCIe “standup” cards</li></ul>
<b>Offline migration to a different VSP</b>	VSPs must be on the same SAN and subnet DIO cannot be used
<b>Serviceguard</b>	VSPs for vPars participating in the same Serviceguard cluster must have the same Fibre Channel and LAN connectivity that physical servers for legacy vPars had

## Review the Virtual Partition command changes

The following changes have been made to the vPar commands. Please review these changes and make any necessary adjustments to your system administration scripts and procedures that are necessary:

- As of v6.1, all commands must be executed from the VSP. No vPar commands are available from the individual vPars. As a result of this, there is no need for secure administration command options for the vPars.
- As of v6.1, only HP-UX 11i v3 virtual partitions are supported. It is possible to maintain HP-UX 11i v2 environments if you are willing to migrate to a standard Integrity VM guest. See the documentation at [hp.com/go/hpux-hpvm-docs](http://hp.com/go/hpux-hpvm-docs).
- All vPartition commands and command flags that directly manage or interact with Extensible Firmware Interface (EFI) and the system firmware have been removed. Examples include `vparenv`, `vparefiutil`, and the options used to manage the EFI variables.
- All command flags related to memory granularity have been removed. Memory is allocated to a vPar in multiples of 64 MB.
- Memory migration, and therefore, the base/floating memory concept, is not supported on v6.1 at this time.
- The Integrity VM commands may also be used to manage the vPars (if desired).
- The Integrity Virtual Server Manager GUI may be used instead of the command line interface (CLI).
- The following list contains all of the supported commands:
  - `vparboot`
  - `vparcreate`
  - `vparmodify`
  - `vparremove`
  - `vparreset`
  - `vparstatus`
  - `vparhwmgmt` (new)—manages the pool of CPU resources dedicated for use by the virtual partitions
  - `vparnet` (new)—creates and manages the virtual network switches (vswitches)
  - `vparconsole` (new)—connects the user to the console of a virtual partition
- The format of the I/O specifications in the commands has changed from hardware path to the Accelerated Virtual I/O (AVIO) format used with Integrity VM:

```
[tape|change|burner|disk|hba|network]:[avio_stor|avio_net]:pcibus:pcislot:aviotgt: \  
{physical device type}:{physical device path}
```

For more information, see the HP vPars and Integrity Virtual Machines V6.1 Administrator Guide available at [hp.com/go/hpux-vpars-docs](http://hp.com/go/hpux-vpars-docs).
- On the Superdome 2, vPars v5 supports operations launched from the Onboard Administrator (OA). When interacting with vPars in this manner, you must specify which nPartition you wish to work in. On vPars v6, all operations are run from the VSP, which is inside of an nPartition, so the options used to specify the nPar ID have been removed.
- In version 6, the vPars are automatically optimized for locality, thus there is neither need nor support for socket local command options and output.
- Command output format has been altered when necessary to adjust for the altered I/O format and capabilities, and the removal of socket local specifications.
- The `getconf` command will no longer return different values for `PARTITION_IDENT` and `MACHINE_IDENT`.

## Plan the new system

### Determine I/O types

This is more than simply counting the number of I/O ports or cards used in the current environment. Since many of the currently shipping I/O cards are significantly faster than those supported by the legacy cell-based servers, and it is possible with vPars v6.1 to allocate the ports on a multi-port card to different vPars, and to actually share physical ports, you probably will not need as many I/O cards on the new system as you have on the old one.

### Determine server type

This is mostly a sizing consideration. The choices include the Integrity based C-class i2 blades, rx2800 i2, and the Superdome 2.

## Extract data from the existing configuration

Using “vparstatus -v”, determine the resources used by the current configuration. Convert these resources to their equivalents for new vPars and record what you need. A few rules to consider are:

- Assume that you should do a 1 for 1 replacement of processor cores in the virtual partitions. Add a minimum of 1 core for the VSP (for vPar management tasks) to get your total requirement.
- Assume that the operating system and key layered software (the “operating environment”) will take approximately 36 GB of disk space. Contact your applications vendors to determine whether your application requirements have changed.
- If your current vPars are running 11i v3, then allocate the same amount of memory to the new ones as you have on the existing ones (unless you are memory limited). If your current vPars are running 11i v2, keep in mind that these will need to be updated to 11i v3, thus you should add at least 1 GB of memory to each of them as the memory requirements for 11i v3 are higher than those for 11i v2.
- Consider how much growth you want to provide for (both in terms of additional vPars and in terms of making the old ones bigger).
- Allocate 20 percent of the total memory for the vPars being migrated (plus any new vPars you plan on creating) to the VSP. For example, if a server will run 6 vPars with a total of 30 gigabytes of memory between them, then allocate an additional 6 gigabytes of memory for the VSP.
- If you are currently running on a cell-based server, assume that you can use NPIV to share Fibre Channel ports without performance impacts (the new Fibre Channel HBAs are much faster than the old ones). Assume that unless you have I/O bound workloads, you can share a Fibre Channel port 8 ways without problem. Bear in mind that if High Availability (HA) is needed, you will have to have paths available over more than one physical port. Also consider the fact that if you use Virtual Connect, the physical HBA may already be divided by the Virtual Connect configuration.
- If you are on a cell-based server, assume that you can use a virtual switch to share a single LAN connection approximately eight ways. Again, please bear in mind any HA requirements you may have. If you need to dedicate a physical LAN port to a single vPar for performance reasons, then you should consider DIO, which will not only dedicate hardware, but will also reduce the load on the VSP.

A sample record of decisions is presented in table 2.

**Table 2.** Example vPar requirements log

vPar name	Cores	RAM	# FC ports (NPIV)	Other non-NPIV storage (file, disk, or lvol backed)		Internal-only virtual LAN	Virtual LANs vswitches	# DIO LANs (physical)
				vPar view	Backing store			
db1	3	6 GB	2	-	-	recovLAN	heartbeat dbLAN mgmtLAN	0
app1	2	4 GB	1	-	-	recovLAN	dbLAN appLAN mgmtLAN	0

app2	2	4 GB	1	DVD	install.iso	recovLAN	dbLAN appLAN mgmtLAN	0
Ignite	1	3 GB	1	-	-	recovLAN	mgmtLAN	0
Spares for future growth	6	12 GB	4	-	-	-	-	2
VSP	1	6 GB	NA	NA	2 disks (mirrored)	-	-	mgmtLAN

## Review the configuration and verify the new system's hardware requirements

While reviewing table 2, you should consider what new server you plan to move to and make sure that:

- The resources added for projected growth will actually fit the planned applications
- Each vPar has connectivity to all of the required LANs and storage
- The resources allocated for the migrating vPars take into account any changed requirements caused by new versions of the applications
- Resources have been allocated for the VSP: at least 1 CPU core, sufficient memory to manage the vPars (~20 percent of the total vPar memory), local storage for the VSP's OS, and networking connections
- The physical server can actually hold all of the planned resources (It is often a good idea to get a server that can expand to twice the capacity initially called for)

## Install the VSP (server or nPartition) with HP-UX vPars and Integrity VM 6.1

Start with the HP-UX 11i v3, March 2012 Update Release (or later). If you purchase either the Virtual Server or Data Center Operating Environments (VSE-OE or DC-OE), then select Integrity vPars and Virtual Machines v 6.1 for installation at the same time as the OS. Select the bundles BB068AA, VirtualBase, and any patches listed in the Integrity vPars and Virtual Machines Release Notes (available at [hp.com/go/vPars](http://hp.com/go/vPars)). If you have purchased the Base OE (BOE) or the HA OE, vPars and Integrity VM version 6.1 is also available as a separately purchased product on the Application Release media. Refer to the product information available at [hp.com/go/vPars](http://hp.com/go/vPars) for more information. If you are migrating to a Superdome 2, configure the nPartitions and stay in nPar mode. Do not go into vPar mode. If you are migrating to c-Class i2 blades, each BladeSystem is a single server and may be administrated in a manner similar to the legacy rackmount servers once physical installation has been completed. The rx2800 is a single server.

Perform the installation:

- Install the HP-UX 11i v3 VSE-OE, explicitly selecting BB068AA and VirtualBase. Refer to the **HP-UX 11i v3 Installation and Update Guide** for the March 2012 Release.
- Apply the recommended configuration changes/performance tunes referenced in the Release Notes.
- Create vswitches for each of the shared and internal LANs desired. Detailed instructions are available in the HP vPars and Integrity Virtual Machines Administrator Guide.
- Configure DIO, if desired. See the Administrator Guide.
- Configure NPIV, in particular, set up the GUIDMGR server if none currently exists, and configure the vPars host to use it. (See the NPIV white paper at [hp.com/go/vPars](http://hp.com/go/vPars)).

# Generate commands for creating the new vPars and infrastructure

In order to create a vPar from table 2, use the `vparcreate` command for each row, with command line parameters based on the table cells as follows:

- **vPar Name:**  
`- p {name}`
- **Cores:**  
`- a core::{number}`
- **RAM:**  
`- a mem::{megabytes}`
- **FC ports: if you have configured your VSP to work with GUIDMGR, then for each of the required virtual HBAs, add**  
`- a hba:avio_stor::npiv:/dev/fcdN`  
 where `/dev/fcdN` is the device file of the physical Fibre Channel port on the VSP that you wish to use with the vPar. If you have decided to manage the virtual World Wide Names (WWNs) manually, then use parameters in the format  
`- a hba:avio:,,{Port WWN},{Node WWN}:npiv:/dev/fcdN`  
 For more information, see HP-UX vPars and Integrity VM NPIV at [hp.com/go/vPars](http://hp.com/go/vPars).
- **Other non-NPIV storage:**  
 For virtual disks and dvds backed through storage mounted to the VSP, the command format is:  
`- a [disk|dvd]:avio_stor::[disk|lv|file]:`  
 where `[disk|dvd]` refers to how the vPar will see the storage and `[disk|lv|file]` refers to how the VSP views the physical backing store.
- **Internal-only virtual LAN and virtual LAN “vswitches”:**  
`- a network:avio_lan::vswitch:{vswitch_name}`  
 Where `{vswitch_name}` is the name of a virtual switch that was created during the VSP creation phase (see **“Install the VSP (server or nPartition) with HP-UX vPars and Integrity VM 6.1”**). The only difference between an internal-only LAN and standard vswitch is whether the vswitch is backed by a LAN port. Each vswitch may be connected to one physical LAN, and each physical LAN may be connected either to a vswitch or be configured for DIO.
- **DIO LANs:** Assuming that you assigned one or more network devices to the DIO resource pool in the **“Install the VSP (server or nPartition) with HP-UX vPars and Integrity VM 6.1”** section, the command parameters would be:  
`- a lan:dio::hwpath:{hardware path}`

## Examples

The following example shows the conversion of a resource planning table entry for a typical vPar with multiple LANs, a virtual DVD, and one virtual HBA:

**Table 3.** Example vPar with multiple LANs

vPar name	Cores	RAM	FC ports (NPIV)	Other non-NPIV storage (file, disk or lvol backed)		Internal-only virtual LAN	Virtual LANs vswitches	# DIO LANs (physical)
				vPar view	Backing store			
my_vPar	2	4 GB	1	DVD	install.iso	recovLAN	dbLAN mgmtLAN	0

**Command:**

```
vparcreate -p my_vPar -a core::2 -a mem::4096 \
-a dvd: avio_stor::file:/images/install.iso

vparmodify -p my_vPar -a hba:avio_stor::npiv:/dev/fcd1

vparmodify -p my_vPar -a network:avio_lan::vswitch:recovLAN

vparmodify -p my_vPar -a network:avio_lan::vswitch:dbLAN \
-a network:avio_lan::vswitch:mgmtLAN
```

The next example covers a vPar with no sharable LANs (other than an internal-only LAN), but with one LAN to be assigned directly to it:

**Table 4.** Example vPar with only an internal LAN

vPar name	Cores	RAM	FC ports (NPIV)	Other non-NPIV storage (file, disk or lvol backed)		Internal-only virtual LAN	Virtual LANs vswitches	# DIO LANs (physical)
				vPar view	Backing store			
vp2	2	4 GB	1	-	-	recovLAN	0	1

**Command:**

```
vparcreate -p vp2 -a core::2 -a mem::4096
vparmodify -p vp2 -a hba:avio_stor::npiv:/dev/fcd1
vparmodify -p vp2 -a network:avio_lan::vswitch:recovLAN
vparmodify -p vp2 -a lan:dio::hwpath:0/0/0/4/0/0/0
```

## Verifying the new vPar and VSP configurations

Prior to attempting to provision the various virtual partitions, confirm that the vPars have the expected resources. This can be done with the `vparstatus`, `vparnet`, and `vparhwmgmt`. Referring back to table 2, if the appropriate commands were run to create these vPars, one would expect the command output to be:

**Figure 1.** Confirming vPar configurations

```
# vparstatus
[Virtual Partition]
Num Name                               RunState      State
=== =====
 1 db1                                  DOWN          Inactive
 2 app1                                  DOWN          Inactive
 3 app2                                  DOWN          Inactive
 4 ignite                                 DOWN          Inactive

[Virtual Partition Resource Summary]
Virtual Partition      CPU           Num   Num   Total MB
Num Name              Min/Max      CPUs  IO    Memory
=== =====
 1 db1                 1/127        3     6     6144
 2 app1                 1/127        2     5     4096
 3 app2                 1/127        2     6     4096
 4 ignite               1/127        1     3     3072

# vparstatus -v
[Virtual Partition Details]
Number:      1
Name:        db1
RunState:    DOWN
State:       Inactive

[CPU Details]
Min/Max:     1/127
System assigned [Count]: 3

[Memory Details]
Total Memory (MB): 6144

[OL* Details]
```

Operation : None  
CPU Count: 0  
Status:

[Authorized Administrators]  
Oper Groups:  
Admin Groups:  
Oper Users:  
Admin Users:

[Remote Console]  
Remote Console Ip Address:  
Remote Console Net Mask:

[IO Details]  
hba:avio\_stor:0,0,0x50014C2000000007,0x50014C2800000007:npiv:/dev/fcd0  
hba:avio\_stor:0,1,0x50014C2000000008,0x50014C2800000008:npiv:/dev/fcd1  
network:avio\_lan:0,2,0xcaffe3509779:vswitch:recovLAN:portid:1  
network:avio\_lan:0,3,0xcaffe350977a:vswitch:heartbeat:portid:1  
network:avio\_lan:0,4,0xcaffe350977b:vswitch:dbLAN:portid:1  
network:avio\_lan:0,5,0xcaffe350977c:vswitch:mgmtLAN:portid:1

[Virtual Partition Details]  
Number: 2  
Name: appl  
RunState: DOWN  
State: Inactive

[CPU Details]  
Min/Max: 1/127  
System assigned [Count]: 2

[Memory Details]  
Total Memory (MB): 4096

[OL\* Details]  
Operation : None  
CPU Count: 0  
Status:

[Authorized Administrators]  
Oper Groups:  
Admin Groups:  
Oper Users:  
Admin Users:

[Remote Console]  
Remote Console Ip Address:  
Remote Console Net Mask:

[IO Details]  
hba:avio\_stor:0,0,0x50014C2000000009,0x50014C2800000009:npiv:/dev/fcd0  
network:avio\_lan:0,2,0x0e1702441afc:vswitch:recovLAN:portid:2  
network:avio\_lan:0,3,0x0e1702441afd:vswitch:appLAN:portid:1  
network:avio\_lan:0,4,0x0e1702441afe:vswitch:dbLAN:portid:2  
network:avio\_lan:0,5,0x0e1702441aff:vswitch:mgmtLAN:portid:2

[Virtual Partition Details]  
Number: 3  
Name: app2  
RunState: DOWN  
State: Inactive

[CPU Details]  
Min/Max: 1/127  
System assigned [Count]: 2

[Memory Details]  
Total Memory (MB): 4096

[OL\* Details]  
Operation : None

CPU Count: 0  
Status:

[Authorized Administrators]  
Oper Groups:  
Admin Groups:  
Oper Users:  
Admin Users:

[Remote Console]  
Remote Console Ip Address:  
Remote Console Net Mask:

[IO Details]  
hba:avio\_stor:0,0,0x50014C200000000a,0x50014C280000000a:npiv:/dev/fcd0  
dvd:avio\_stor:1,0,0:file:/images/install.iso  
network:avio\_lan:0,2,0xa68285ade9cf:vswitch:recovLAN:portid:3  
network:avio\_lan:0,3,0x0e1702441ad1:vswitch:appLAN:portid:2  
network:avio\_lan:0,4,0x0e1702441ad2:vswitch:dbLAN:portid:3  
network:avio\_lan:0,5,0x0e1702441ad3:vswitch:mgmtLAN:portid:3

[Virtual Partition Details]  
Number: 4  
Name: ignite  
RunState: DOWN  
State: Inactive

[CPU Details]  
Min/Max: 1/127  
System assigned [Count]: 1

[Memory Details]  
Total Memory(MB): 3072

[OL\* Details]  
Operation : None  
CPU Count: 0  
Status:

[Authorized Administrators]  
Oper Groups:  
Admin Groups:  
Oper Users:  
Admin Users:

[Remote Console]  
Remote Console Ip Address:  
Remote Console Net Mask:

[IO Details]  
hba:avio\_stor:0,0,0x50014C200000000b,0x50014C280000000b:npiv:/dev/fcd1  
network:avio\_lan:0,2,0x0e1702441afc:vswitch:recovLAN:portid:4  
network:avio\_lan:0,5,0x0e1702441aff:vswitch:mgmtLAN:portid:4

```
# vparnet
```

Name	Number	State	Mode	NamePPA	MAC Address	IPv4 Address
recovLAN	1	Up	Shared		N/A	N/A
heartbeat	2	Up	Shared	lan47	0x0022640fd474	
dbLAN	3	Up	Shared	lan16	0xf4ce46fdfe78	
appLAN	4	Up	Shared	lan17	0xf4ce46fdfe7c	
mgmtLAN	5	Up	Shared	lan38	0xf4ce46fd0f90	

```
# vparhwmgmt -p cpu -l  
Cores allocated for vPars or VMs: 15  
Cores allocated for VSP: 1  
Cores not usable: 0  
Total number of operable system cores: 16  
Total number of system cores: 16  
Core pool OL* policy: GUEST
```

The example contains the summary vPar configuration output (`vparstatus`) which is useful for sanity checks:

- the verbose output (`vparstatus -v`) verifies that the I/O has been properly configured
- the `vparnet` output makes sure all of the virtual LANs are as expected
- the `vparhwmgmt` output shows how many resources are reserved for use by the VSP.

## Provision virtual partitions

Once the VSP is up and the virtual partitions have been created, provision each vPar with an OE, applications, and data. Given the assumption that most of your vPars will actually be migrated from legacy hardware, decide whether to create fresh boot disks and attach the old data, or to keep your existing boot images and update them.

### Starting with fresh boot images

The recommended method, particularly if you are migrating from a PA-RISC system or an older version of HP-UX, would be to start with a new LUN and to use Ignite-UX to do a fresh install of HP-UX (this requires an AVIO NIC attached to a LAN based vswitch or an internal only LAN with a running vPar functioning as an Ignite server). This is the fastest, easiest, and least error prone method of getting an OS image. It has the added advantage of not carrying forward any outdated HP or third-party software, such as the old vPars product.

If you choose this approach, be sure that you are installing the March 2012 update for HP-UX 11i v3 or later. The actual operating environment you use does not matter from the virtualization standpoint, but other considerations, such as high availability, may apply. See the **HP-UX 11i v3 Installation and Update Guide**. Once the root disk is installed, it will be possible to either mount the old data and application disks, or to mount fresh volumes and copy the data.

### Migrating disks from legacy servers

**Please note that this is not recommended.**

While it is possible to update a sufficiently large root disk for HP-UX, in many cases, between file system growth and changes to kernel tunable parameters, it is often faster and cleaner to start with a new disk.

If you still wish to migrate your boot disk, HP does support creating boot disks on legacy HP Integrity servers and using them to boot the new virtual partitions. This can be achieved in at least three ways:

1. Create a dynamic root disk (DRD) clone and update the clone to the HP-UX March 2012 release or later
2. Create an Ignite Recovery image and depots with the additional required software
3. Update the boot disk in place and then move it (or if it is on the SAN, remount it)

This section provides a summary of the document “Migrating an Integrity HP-UX 11i v3 instance to new hardware.” For detailed coverage of this topic, please refer to the full paper which may be found online at [hp.com/go/DRD](http://hp.com/go/DRD) under “Learn more—White paper: Migrating an Integrity HP-UX 11i v3 instance to new hardware.”

The following table makes some suggestions in this area.

**Table 5.** Update method selection criteria

Situation	Consider using this method
If you are moving the entire system—data as well as boot disks	Move the boot disk or DRD clone, along with all data. Device files will be preserved and very little additional setup will be needed.
If you do not have an Ignite-UX server set up already	Move the boot disk or a DRD clone.
If you have many boot disks on a given system or hard partition, thus making it challenging to identify a particular boot disk	Deploy an Ignite-UX recovery image.
If your root disk is managed by VxVM	Deploy an Ignite-UX Recovery image. DRD rehosting does not support VxVM roots.
If you only use internal boot disks that are not hot-pluggable	Deploy an Ignite-UX recovery image.
If you want to move an entire set of vPars	Set up the vPars on the VSP, then either move DRD clones or deploy Ignite-UX recovery images.

## Validating that the old disks can be updated and reused

- Integrity server disks only, not PA-RISC servers
- Supported storage
- 11i v3 to start with

## Overview of the DRD method

The overall approach used here is to create a DRD clone on an older “source” system, modify it to support the new “target” system model, and move it to the new target system. The recommended steps will detail the upgrades and changes that must be accomplished on the source system, as well as those that must be defined for automatic consumption during boot of the target system.

### Note:

A system administrator may choose to move the actual boot disk from the source system to the target. To do this, the following changes should be applied to the boot disk rather than to the DRD clone. However, this makes it somewhat more difficult to revert to the original hardware if issues are encountered. For this reason, the following procedure is described for a DRD clone.

## Assumptions for the DRD method

- DynRootDisk is installed on the source system and on the DRD clone of the source system at release B.1131.A.3.6 or above. This release, supplied with the September 2010 media, includes support for kernel management on the clone using “drd runcmd mk\_kernel”, “drd runcmd kcmodule” and “drd runcmd kconfig” (as well as the previously supported “drd runcmd ktune”).
- The release of HP-UX on the source system is 11i v3.
- The root volume group is managed by LVM.

## Summary of steps for the DRD method

In some cases, the goal is to move a pre-existing HP-UX instance to new hardware; in other cases, the goal is to deploy a very similar (same HP-UX release, same patches, etc.) system with a different network identity (hostname, MAC, and IP addresses, etc.). The difference between the steps needed for these two scenarios is small, so both scenarios are covered by these steps. The following steps will be used in both scenarios:

1. Create a DRD clone of the source system on storage that can be moved to the target.
2. Modify file system sizes on the clone if needed.
3. Identify and install additional target software on the clone (see the system Errata document for the new vPar machine. Errata documentation may be found by a Web search on the system model and the string “Errata”<sup>1</sup>).
4. Determine additional kernel content that is needed on the target (see Step 4 of the “Use the DRD clone to migrate to new hardware” section of the white paper **Migrating an Integrity HP-UX 11i v3 instance to new hardware**).
5. Build a kernel on the clone suitable for the target (see step 5 of the “Use the DRD clone to migrate to new hardware” section of the white paper **Migrating an Integrity HP-UX 11i v3 instance to new hardware**).
6. Optional: Adjust target kernel tunable parameters.
7. Set the system identity on the clone for boot on the target.
8. Mark the clone LUN for identification from the target EFI.
9. Move scenario: Disable the source system.
10. Move storage.
11. Boot the clone on the target system.
12. Test the target system.
13. If the target does not satisfy expectations, repeat the process, making any necessary adjustments.

---

<sup>1</sup> Note: any software downloaded from the Web must be converted from a depot file to a “standard” depot by using the swcopy command. DRD does not support running swinstall on a depot file to a clone disk.

## Overview of the Ignite-UX recovery method

The overall approach used here is to deploy an Ignite-UX recovery archive of the HP-UX instance running on older hardware, called the source system, to the new system, called the target. The recommended steps detail the upgrades and changes that must be accomplished on the source system prior to creating the recovery archive, on the Ignite-UX server after the recovery archive is created, and during boot of the target system.

In some cases, the goal is to move a pre-existing HP-UX instance to new hardware; in other cases, the goal is to deploy a very similar system—same HP-UX release, same patches, etc.—with a different network identity—hostname, MAC, and IP addresses, etc. The difference between the steps needed for these two scenarios is small, so both scenarios are covered here. To distinguish these similar scenarios, the first is called the move scenario, and the second is called the clone scenario.

## Assumptions for the Ignite-UX recovery method

It is assumed that the HP-UX major release installed on the source system is 11i v3. The use of the “match” specification when installing an Operating Environment requires 11i v3.

## Summary of steps for the Ignite-UX recovery method

1. Install the latest IGNITE bundle on the Ignite-UX server.
2. Make a recovery archive of the source system.
3. Copy the source client configuration to the target configuration.
4. Clone scenario: Give the target read access to the source recovery archive.
5. Remove or modify the network information for the target (see Step 5 of the “Detailed Migration Steps for the Ignite-UX Recovery Method” chapter of the white paper **Migrating an Integrity HP-UX 11i v3 instance to new hardware**).
6. Change a recovery-related variable(s) in `control_cfg` (see Step 6 of the “Detailed Migration Steps for the Ignite-UX Recovery Method” chapter of the white paper **Migrating an Integrity HP-UX 11i v3 instance to new hardware**).
7. Create depots with additional software for the target (see the HP-UX Errata for your target server).
8. Ensure `iux_postload` scripts from the recovery archive run correctly.
9. Create config files for the depots.
10. Add the new config files to the target CINDEX.
11. Modify file system sizes if needed.
12. Check configuration syntax.
13. Move scenario: Disable the source system.
14. Deploy the configuration to the target system.
15. Boot and test the target system.
16. If the target is not as desired, return to the source system, and make adjustments as necessary.

## Post migration steps

### Confirm successful vPar boot and data visibility

You can boot and install a virtual partition in several ways. If you have assigned pre-installed boot disks, a physical DVD drive with the installation disk in it, or an ISO image of an installation disk, then you simply enter:

```
# vparboot -p {vpar_name}
# vparconsole -p {vpar_name}
```

and interact with the EFI prompts in the same way you would on a standalone system:

```
Shell> fs0:
fs0:\> hpux
```

If you wish to boot from an Ignite-UX server, then boot the vPar to EFI and create a dbprofile with the following command:

```
Shell> dbprofile -dn mybootprofile -sip {IP_address_of_ignite-server}  
-cip {IP_address_of_vPar} -gip {IP_address_of_gateway} -m {network_mask}  
-b "/opt/ignite/boot/nbp.efi"
```

Once you have done that, use the dbprofile to boot the install kernel:

```
Shell> lanboot select -dn mybootprofile
```

Once you have installed HP-UX (11i v3, March 2012 update release or later), boot the system to a regular shell prompt and ensure that you can see all of the disks and devices that you expect.

## Update applications and I/O paths as necessary

HP recommends that you update your applications to versions that explicitly support HP-UX 11i v3. Also, if the applications that you are using store hardware paths in configuration files, you will have to edit these files to reflect the new paths. If that is necessary, HP recommends that you switch to using the new “agile” addressing model so that future edits to configuration files will be unnecessary.

## Conclusion

With the release of HP-UX vPars v6.1, take advantage of the new, faster, and more efficient servers from HP and the benefits of the mission-critical Converged Infrastructure while keeping the cost savings of virtualization. Consider whether it makes sense to migrate your workloads to the new bladed architectures.

## For more information

Get a more in-depth overview of Virtual Partitions version 6.1 by visiting [hp.com/go/vPars](http://hp.com/go/vPars).

---

### Get connected

[hp.com/go/getconnected](http://hp.com/go/getconnected)

Current HP driver, support, and security alerts  
delivered directly to your desktop

© Copyright 2012 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.

4AA4-2782ENW, Created August 2012

