HP-UX vPars and Integrity VM V6.1 Release Notes

HP Part Number: 5900-1578
Published: March 2012, Edition 1.0

Legal Notices

Acknowledgments

Revision History

Manufacturing Part Number	Supported Operating Systems	Supported Versions	Document Edition Number	Publication Date
5900–1578	HP-UX	11i v3	1.0	March 2012

Contents

Tables

Examples

1 Introduction

Thank you for installing the vPars and Integrity VM V6.1 bundles (Virtualization Services Platform (VSP) software bundle, BB068AA, and vPars/VM software bundle, VirtualBase). This chapter introduces the release notes for the "converged" vPars and Integrity VM V6.1 product.

Install the HP-UX GUID Manager software, which is included with the HP-UX 11i v3 March 2012 release, or can be obtained by going to http://software.hp.com and searching for the HP-UX GUID Manager.

The following patches are required:

- On the VSP (formerly known as the VM Host):
  - PHSS_42795 1.0 HPVM B.06.10 CORE PATCH
  - PHSS_42796 1.0 HPVM B.06.10 VMAGENT
  - PHSS_42797 1.0 HPVM B.06.10 AVIO-HVSD
- On the vPar/VM:
  - PHSS_42798 1.0 HPVM B.06.10 AVIO-GVSD

HP recommends that you also install the following patch on the VSP:

- PHSS_42798 1.0 HPVM B.06.10 AVIO-GVSD

You can obtain these patches from the HP Support Center website at HP Support Center.

NOTE: Future updates to the vPars and Integrity VM V6.1 product will be released as vPars and Integrity VM product patches.

AVIO drivers are now part of the vPars and Integrity VM product bundle.

HP recommends installing only the AVIO drivers that are included with HP-UX 11i v3 March 2012 release with this version of vPars and Integrity VM V6.1.

This manual is organized in the following way:

- This introductory chapter contains a list of the new features provided in the new release, as well as a description of the manual organization. It also may contain information of general interest to the release and documentation issues.
- Chapters 2 and 3 provide release notes, relevant to the V6.1 release.
- Chapter 11 provides the support policy tables for vPars and Integrity VM V6.1 and its components.

For the most up-to-date information about HP-UX vPars and Integrity VM V6.1, see the documentation on the Business Support Center website:

http://www.hp.com/go/virtualization-manuals

For the most recent information relevant to this update, see the `README.txt` file at the following location:

`/opt/hpvm/doc/README.txt`

1.1 Features and Enhancements in vPars and Integrity VM

The features in the following list have been included in this release of vPars and Integrity VM V6.1:

- VSP and vPars/VMs support:

    - HP-UX 11i v3 VSP — HP-UX 11i v3 March 2012 release
    - For vPars: HP-UX 11i v3 September 2011 and March 2012 guests (Intel® Itanium® Processor 9300 series VSPs only)
    - For VMs: HP-UX 11i v3 March 2010 through March 2012 guests
- Support for vPars and Integrity VM functionality in a single, common, easy-to-use management environment.
- Support for the direct I/O (DIO) networking functionality.
- Support for Matrix OE 7.0 fluid cross-technology logical servers with virtual machines. For additional details, see the *HP Matrix Operating Environment 7.0 Logical Server Management User Guide*.

1.2 Version Support

The following table lists support in HP-UX vPars and Integrity VM V6.1 for the various HP Integrity servers and HP Integrity server blades:

Table 1 vPar, VM and VSP Support

Virtual Partitions	Virtual Machines	VSP on HP Integrity server blades[1]	VSP on HP Integrity Superdome 2[1]	VSP on HP Integrity rx2800 i2 Servers[1]	VSP on HP Integrity Servers (9000 Series)
N.A.	HP-UX 11i v3 1003	Yes	No	No	Yes
N.A	HP-UX 11i v3 1009	Yes	Yes	Yes	Yes
N.A.	HP-UX 11i v3 1103	Yes	Yes	Yes	Yes
HP-UX 11i v3 1109	HP-UX 11i v3 1109	Yes	Yes	Yes	Yes
HP-UX 11i v3	HP-UX 11i v3	Yes	Yes	Yes	Integrity VM — Yes vPars — No

[1] Support for the latest Intel® Itanium® Processor 9300 Series

1.3 Changes in Support

The following list contains features that are not supported in V6.1:

- Use of the `hpvmupgrade` and the `pv2assist` commands.
- Use of VIO.

  In addition, no new I/O hardware, such as FCoE CNA's is supported with VIO. To determine if your VMs are using VIO, run `hpvmstatus -d` for all your guests. Look for storage devices that use `scsi` and look for network devices that use `lan`:

  ```
  # hpvmstatus -P guestname  -d | grep -w scsi
  # hpvmstatus -P guestname  -d | grep -w lan
  ```

  A VLAN tag can be configured on an AVIO localnet vswitch port. The guest lan bound to this vswitch port can be used only to communicate with another guest lan bound to a vswitch port that does not have a VLAN tag configured.

- Microsoft Windows, Linux , and OpenVMS guests are not supported. Version 4.2.5 was the last version to support Microsoft Windows and Linux guests.
- Use of legacy device special files (DSFs) to define virtual storage.

  Support for the use of legacy DSFs to define virtual storage (including virtual disks and DVDs) are deprecated. Customers should begin planning to use persistent (agile) DSFs when defining their virtual storage devices. HP recommends the use of persistent DSFs (for example, those with pathnames such as `/dev/rdisk/disk##`) for better storage availability and reliability.

  To check for the use of legacy DSFs, use the following command:

  ```
  # hpvmstatus -P guestname -d | grep -w rdsk
  ```

2 Installation Notes

This chapter contains notes about installing and upgrading HP-UX vPars and Integrity VM V6.1 and associated software on the VSP system.

2.1 Installing HP-UX vPars and Integrity VM 6.1

This section describes information about installing the HP-UX vPars and Integrity VM 6.1 product bundle and associated software on the VSP system.

HP-UX vPars and Integrity VM 6.1 is supported on HP Integrity servers or nPartitions running HP-UX 11i v3 March 2012 release. When you upgrade or reinstall vPars or Integrity VM, guests are stopped, but they are not removed. When the new version of HP-UX vPars and Integrity VM starts, the virtual machines or vPars might also start, depending on the setting of the boot attribute.

NOTE: The VSP system is not a general-purpose system; it is dedicated to the hosting of virtual machines and virtual partitions. After you install vPars and Integrity VM, no operating system or process management reconfiguration should be performed outside those provided by the vPars and Integrity VM interfaces. Specific examples of actions that are not supported on the VSP system include:

- Changing priorities or scheduling attributes of processes on the VSP system.
- Modifying kernel tunables in any way.

NOTE: Do not shut down or reboot the VSP while any vPars or VMs are running.

For complete information about the requirements for installing vPars and Integrity VM, see the *HP-UX vPars and Integrity VM 6.1 Administrator Guide* at http://www.hp.com/go/hpux-hpvm-docs.

2.2 Upgrading to vPars and Integrity VM V6.1

You can upgrade from HP-UX Virtual Partitions V6.0 to vPars and Integrity VM V6.1 and from Integrity VM V3.5, V4.1, V4.2, V4.2.5, and V4.3. Upgrading from Integrity VM V4.0 is not supported. Note, only HP-UX 11i v3 guests are supported in V6.1.

3 Issues in this release

This chapter contains notes about creating and configuring virtual machines on the VSP system.

3.1 Changes and Issues in this Release

This section describes new information and issues in this release.

3.1.1 hpvmhwmgmt might add ports in link aggregates into the DIO pool

The hpvmhwmgmt command might put a port used by a link aggregate into the direct I/O (DIO) networking pool. A LAN interface cannot be used by both APA and DIO at the same time. Before adding a LAN interface into the DIO pool, make sure that it is not being used by an APA interface.

To find out whether a DIO capable lan device is being used in a link aggregate, use the ioscan and nwmgr commands. For example, add 0/0/0/9/0/0/0 into the DIO pool, a DLA device with 0/0/0/9/0/0/1 as its other sibling, and check whether both of these are being used by a link aggregate:

1. Use ioscan to find the instance numbers of these lan interfaces:

```
# ioscan -kfnC lan

lan      16  0/0/0/9/0/0/0  iexgbe   CLAIMED     INTERFACE    HP 454521-002E PCIe 2-p …
lan      17  0/0/0/9/0/0/1  iexgbe   CLAIMED     INTERFACE    HP 454521-002E PCIe 2-p …
```

   Now the interfaces are PPAs 16 and 17.

2. Use nwmgr to see whether lan16 and lan17 are being used by a link aggregate. If you see lan9XX in the "Related Interface" column for the lan interface in which you are interested, then it is being used by APA:

```
# nwmgr
Name/          Interface Station          Sub-   Interface      Related
ClassInstance  State     Address         system   Type           Interface
============== ========= ============== ======== ============== =========
lan0           UP        0x78ACC089BF02 iexgbe   10GBASE-KR
          …
lan7           DOWN      0x78ACC089BF09 iexgbe   10GBASE-KR
lan16          UP        0x68B599B54220 iexgbe   10GBASE-KR     lan901
lan17          DOWN      0x68B599B54224 iexgbe   10GBASE-KR
         ..
lan900         DOWN      0x000000000000 hp_apa   hp_apa
lan901         UP        0x68B599B54220 hp_apa   hp_apa
lan902         DOWN      0x000000000000 hp_apa   hp_apa
lan903         DOWN      0x000000000000 hp_apa   hp_apa
lan904         DOWN      0x000000000000 hp_apa   hp_apa
#
```

   In this case, lan16 is being used by APA interface lan901, but not lan17.

3.1.2 Analyzing failures when DIO resources in use by the VSP

If hpvmhwmgmt -p dio -a hwpath fails due to a resource being in use by the VSP, check the /var/adm/cra.log file for additional information on the resources in use by the VSP. The following example shows the type of error you might see in this case:

```
system# hpvmhwmgmt -p dio -a 0/0/0/4/0/0/0
hpvmhwmgmt: ERROR - Resources for: '0/0/0/4/0/0/0' are in use by your host.
hpvmhwmgmt: ERROR - could not reserve hwpath: '0/0/0/4/0/0/0' for DIO.
hpvmhwmgmt: ERROR - failed to reserve and store hw path: '0/0/0/4/0/0/0' for DIO.
hpvmhwmgmt: Unable to manage dio pool resource.
```

3.1.3 DIO management device file appears stale to lssf

The `/dev/hpvmdio` device file is listed by `lssf -s` as a stale device file, however, it is not a stale device file and should not be removed from the VSP system. Removing the `/dev/hpvmdio` device will result in no longer being able to see the DIO management pool with the `hpvmhwmgmt -p dio` commands. If the `/dev/hpvmdio` character device file is removed, use the `mknod` command to recreate it. The `mknod` command requires both a major and minor number. The major number can be found by running `lsdev` and locating the character number for the *`hpvmdio`* driver. The minor number is 0xffffff.

3.1.4 DIO Broadcom stops receiving link notifications

If a DIO Broadcom link is kept down and a vPar or VM is rebooted, the DIO interface will not receive further link notifications. To prevent this, bring up the DIO link before booting the vPar/VM; otherwise, the only way to recover from this situation to reboot the VSP.

3.1.5 DIO Broadcom does not see bandwidth changes

DIO Broadcom will see only dynamic bandwidth changes when all functions of the device are owned by a vPar/VM or still owned by the VSP. If only some of the device's functions are assigned to the vPar/VM, bandwidth changes will take effect after VSP reboot.

To workaround this situation, all the interfaces from a device must be assigned to a vPar/VM or VSP, and nothing should be left in the DIO pool.

3.1.6 DIO-capable functions might become inconsistent with information in vPar/VM device database

Under certain conditions, the ownership of DIO-capable functions might become inconsistent with the information in the vPar/VM device database such that `hpvmhwmgmt -p dio -l` indicates that a DIO-capable function is owned by `host`. However, the `hpvmdevmgmt -l gdev` command indicates that it is assigned to the vPar/VM resource pool as a `gdev`.

When this happens, do the following:

1. Save your vPar or VM configurations that specify how you are using DIO. For example:

   ```
   # hpvmhwmgmt -p dio -l > /tmp/save_diopool_info
   # hpvmstatus -P  name -d | grep dio > /tmp/dioconfig_name
   ```

   Where: *`name`* is the name of each vPar or VM that is configured with DIO.

2. De-assign the DIO resources. To do so, you must use `hpvmmodify -P name -d resource` to delete the resource from any vPar or VM to which it might be assigned, then delete the resource from the device database with the following procedure:

   ```
   # hpvmdevmgmt -m gdev:hwpath:attr:NODELETE=NO
   # hpvmdevmgmt -d gdev:hwpath
   ```

   If you want to use the DIO-capable function again as an vPar/VM resource, at this point you can re-add it to the DIO resource pool using the following command:

   ```
   # hpvmhwmgmt -p dio -a hwpath
   ```

   Then re-add the resource to any vPar or VM.

3.1.7 DIO limitations

DIO has the following limitations:

- Boot over a DIO device is not supported. Install or boot a vPar/VM over LAN must use AVIO devices.
- No error recovery on HP Integrity Superdome 2 systems. Error recovery will not happen on DIO devices when a recoverable error occurs. The guest will be punished.

- Cannot dynamically add, delete, modify DIO functions online (that is, virtual OLRAD); such `hpvmmodify` operations have to be done when the vPar/VM is down.
- Removing `hpvmdio`-owned devices while the VSP is down might result in broken VM configurations that must be fixed manually.
- Cannot dynamically add, remove vPar/VM memory with DIO functions assigned.
- No Online VM Migration support.
- No suspend and resume support for vPars/VMs configured with DIO. These operations fail if the vPar/VM has a DIO device.
- No firmware upgrades to DIO devices inside the vPar/VM.
- No more than 16 DIO functions can be assigned to a single vPar/VM
- No support for HP Auto Port Aggregation (APA).

3.1.8 No warnings when attempting to change memory size of a running vPar

It is not possible to change the memory size of a running vPar. However; the `hpvmmodify` command and the `vparmodify` command do not print any warnings or errors if this is attempted.

3.1.9 The `hpvmmodify` command might fail when it should succeed

The `hpvmodify` command invoked on a running VM might fail when it should succeed. When failing, the following error message is displayed:

```
resource intent failed 'Can't get the resource maxima.'
vPar/VM vm_name configuration problems:
    Error 1: Internal error -1 when attempting to use the ragent 'intent' interface
hpvmmodify: Unable to modify the vPar or VM.
```

This failure might happen only when the same processor is used by several VMs.

3.1.10 The `hpvmmodify` command might succeed when it should fail

When the `hpvmmodify` command is invoked on a running VM to increase its minimum entitlement (`-e` option), it might succeed when it should fail. In this case, the `hpvmmodify` command displays the following warning message:

```
# hpvmmodify -P vm_name -e minimum entitlement
vPar/VM vm_name configuration problems:
    Warning 1: vPar/VM CPUs (n) greater than available limit (0).
These problems may prevent the vPar or VM vm_name from starting.
hpvmmodify: The modification process is continuing.
hpvmmodify: The modification process is continuing.
```

As a consequence, a reserved VM might not be able to start, because not enough CPU is available. Another consequence is that the VSP resources might become overcommitted:

```
# hpvmstatus -S
                [HPVM Server System Resources]

        *** VSP resources are over-committed ***
        vPar/VM types supported by this VSP = vPar, Shared
        Processor speed = 1595 Mhz
                                .
                                .
```

To work around this overcommitted condition, reduce the minimum entitlement of the VM whose entitlement increase created the condition.

3.1.11 Converting VMs to vPars causes LAN IDs to shift

The I/O hardware paths on VMs do not have a cell number in the path. For vPars, I/O resources are shown under a cell, and because of that, I/O hardware paths have an additional "0/" prefix.

When a VM is converted to a vPar, the vNIC hardware path will be different, so `ioscan` enumerates it with a different PPA (lan index).

Workaround

Boot the vPar to single user mode, fix the `/etc/rc.config.d/netconf` file with the new PPA number, and then switch to run level 3.

3.1.12 SMH is unresponsive when running stress tests on a VSP

Applications like SMH (which needs significant CPU bandwidth) are not likely to be very responsive when the VSP cores are already under heavy load servicing vPar/VM I/O requests. HP recommends that you increase the number of VSP CPU-cores under such circumstances.

3.1.13 WARNING: VCPU0 not scheduled message can be ignored

In vPars V6.1, messages similar to the following are occasionally seen during the initial boot of a vPar:

```
WARNING: VCPU0 not scheduled for NNNNN ms" messages in hpvm_mon_log
```

They can occur during the early portion of booting the vPar before HP-UX is launched into the vPar, when either the EFI layer or the boot loader is running.

You can safely ignore these messages.

3.1.14 VSP requires dedicated CPU to support a vPars

If you plan to use vPars, you need to setup the VSP with a dedicated CPU to support them. Virtual machines do not require a dedicated CPU.

3.1.15 Use of `-F` with `hpvmmigrate` on a suspended VM can cause VM to be not runnable on both source and target

You can use the `-F` to override resource checking during a VM migrate operation. However, if you use this option on a suspended VM, and you have also used NPIV on the VM and a target VSP that cannot support that NPIV port, the VM migrate operation will fail, even with the `-F` option. This will cause the VM to be nonrunnable on both the source VSP and target VSP.

To re-enable the VM on the source VSP, remove the suspend state of the VM on the source VSP with the following commands:

```
hpvmmodify -P vmname -x suspend_file=delete –F
hpvmmodify -P vmname -x register_status=enabled
```

3.1.16 Integrity VM and vPar CLI commands experience poor performance when there are numerous devices on the VSP

Commands like `vparmodify`, `hpvmmodify`, `hpvmcreate` and `hpvmclone`, (commands used to modify the vPar or VM configuration), experience slow performance when there are numerous devices available on the VSP, or configured in the vPar and/or VM configurations. When you have a large number of devices, it is more than likely that the majority of those devices are storage devices. If storage devices are being exposed to the VSP from a SAN and then individually mapped to vPar/VM configurations, alternatively, you can map SAN-based LUNs directly to the VMs or vPars using NPIV. Replacing individually mapped SAN-based LUNs with one or more virtual HBAs using NPIV ports, reduces the number of devices that need to be managed, and thus improves the CLI performance. For information about using NPIV with vPars and Integrity VM, see the *HP-UX vPars and Integrity VM V6.1 Administrator Guide* on the BSC website: http://www.hp.com/go/virtualization-manuals.

3.1.17 Adding NPIV HBAs to existing guest with AVIO storage HBAs might cause a problem

When adding NPIV HBAs to an existing guest that has AVIO storage HBAs, the next boot of the guest may panic with the message "`panic: GVSD: Unable to allocate iocb during initialization`", due to the known defect.

To work around this problem, reduce the number of AVIO storage HBAs by consolidating LUNs into fewer AVIO storage HBAs. AVIO allows up to 128 AVIO LUNs under 1 AVIO HBA. The following example shows how to consolidate 20 AVIO LUNs from 20 AVIO HBAs to 1 AVIO HBA:

1. Check status, device info, and perform an `ioscan`. Notice that the example begins with 20 AVIO LUNs under 20 AVIO HBAs:

```
# hpvmstatus -d -P guest
disk:avio_stor:1,3,0:disk:/dev/rdisk/disk6199
disk:avio_stor:1,4,0:disk:/dev/rdisk/disk6520
…
disk:avio_stor:7,3,6:disk:/dev/rdisk/disk6201
disk:avio_stor:7,4,6:disk:/dev/rdisk/disk6413

# hpvmdevinfo
Device Type  Bus,Device,Target  Host Device Name     Virtual Machine Device Name
===========  =================  ================     ===========================
disk         [1,3,0]            /dev/rdisk/disk6199  /dev/rdisk/disk263
disk         [1,4,0]            /dev/rdisk/disk6520  /dev/rdisk/disk266
…
disk         [7,3,6]            /dev/rdisk/disk6201  /dev/rdisk/disk326
disk         [7,4,6]            /dev/rdisk/disk6413  /dev/rdisk/disk333

# ioscan -kfNd gvsd
Class     I  H/W Path  Driver S/W State   H/W Type     Description
==========================================================
ext_bus   7  0/1/3/0   gvsd   CLAIMED     INTERFACE    HPVM AVIO Stor Adapter
ext_bus   8  0/1/4/0   gvsd   CLAIMED     INTERFACE    HPVM AVIO Stor Adapter
…
ext_bus  36  0/7/3/0   gvsd   CLAIMED     INTERFACE    HPVM AVIO Stor Adapter
ext_bus  37  0/7/4/0   gvsd   CLAIMED     INTERFACE    HPVM AVIO Stor Adapter
```

2. Halt the guest to change the configuration.
3. Run `hpvmmodify` to delete the 20 AVIO LUNs from the current configuration:

```
# hpvmmodify -P guest -d disk:avio_stor:1,3,0:disk:/dev/rdisk/disk6199
# hpvmmodify -P guest -d disk:avio_stor:1,4,0:disk:/dev/rdisk/disk6520
...
# hpvmmodify -P guest -d disk:avio_stor:7,3,6:disk:/dev/rdisk/disk6201
# hpvmmodify -P guest -d disk:avio_stor:7,4,6:disk:/dev/rdisk/disk6413
```

4. Run `hpvmmodify` to recreate the 20 AVIO LUNs under a single AVIO HBA:

```
# hpvmmodify -P guest -a disk:avio_stor:2,0,0:disk:/dev/rdisk/disk6199
# hpvmmodify -P guest -a disk:avio_stor:2,0,1:disk:/dev/rdisk/disk6520
...
# hpvmmodify -P guest -a disk:avio_stor:2,0,18:disk:/dev/rdisk/disk6201
# hpvmmodify -P guest -a disk:avio_stor:2,0,19:disk:/dev/rdisk/disk6413
```

5. Restart the guest. Notice that 20 AVIO LUNs are now under 1 AVIO HBA.

```
# hpvmstatus -d -P guest
disk:avio_stor:2,0,0:disk:/dev/rdisk/disk6199
disk:avio_stor:2,0,1:disk:/dev/rdisk/disk6520
…
disk:avio_stor:2,0,18:disk:/dev/rdisk/disk6201
disk:avio_stor:2,0,19:disk:/dev/rdisk/disk6413

# hpvmdevinfo
Device Type  Bus,Device,Target  Host Device Name     Virtual Machine Device Name
===========  =================  ================     ===========================
disk         [2,0,0]            /dev/rdisk/disk6199  /dev/rdisk/disk263
disk         [2,0,1]            /dev/rdisk/disk6520  /dev/rdisk/disk266
…
```

```
disk         [2,0,18]           /dev/rdisk/disk6201  /dev/rdisk/disk326
disk         [2,0,19]           /dev/rdisk/disk6413  /dev/rdisk/disk333

# ioscan -kfNd gvsd
Class     I  H/W Path  Driver S/W State   H/W Type     Description
===========================================================
ext_bus  12  0/2/0/0   gvsd   CLAIMED     INTERFACE    HPVM AVIO Stor Adapter
```

3.1.18 Simultaneous interrupt migration operations in both VSP and VM might cause VSP system crash

In some scenarios, performing simultaneous interrupt migration operations in both the VSP and the VM might result in a VSP system crash.

To work around this, HP recommends that you do not perform interrupt migration operations simultaneously on both the VSP and the VM. There is no harm in performing interrupt migration operations individually on the VSP and the VM.

3.1.19 Copy of a vPar/VM might be left in runnable state if migration fails

If a migration fails, a vPar/VM might be left in a runnable state (but not running) on the target VSP. If this happens, further attempts at migration will fail with a message that the vPar/VM exists and is runnable on the target VSP.

To work around this, either remove the vPar/VM with the `hpvmremove` command, or set the vPar/VM not runnable using the `hpvmmodify -x runnable_status=disabled` command.

3.1.20 Changes to the Default EFI AVIO Storage Driver Enumeration Policy

Prior to HP Integrity Virtual Machines V4.3, the default policy of the EFI AVIO storage driver was to enumerate and configure all AVIO storage devices defined in a guest's configuration. This would make all AVIO storage devices visible to EFI. With V4.3 (and later), the default enumeration policy is to enumerate and configure only those AVIO storage devices that are present as EFI boot options. The enumeration policy is also differentiated by storage type: SCSI (non-NPIV) LUNs, and FC (NPIV) LUNs.

The change to the enumeration policy has been made to reduce the time to boot to EFI in the cases of guest configurations with large amounts of AVIO storage devices configured. This policy change affects only the devices enumerated or configured at the EFI level, not at the guest operating system level.

As a result of the default enumeration policy change, some attempted operations within EFI (such as boot management, or new guest installations) may fail because of non-present AVIO storage devices. To make the AVIO storage devices visible to EFI, the enumeration policy can be changed to enumerate or configure all AVIO storage devices.

Use the drvcfg EFI utility to change the enumeration policy to do the following:

- Enumerate boot LUNs only. (Default policy)
- Enumerate all LUNs.

You can set the enumeration policy separately for SCSI (non-NPIV) LUNs and FC (NPIV) LUNs; however, setting the policy to enumerate all LUNs (especially FC LUNs) might result in long guest boot times in configurations with a large number of LUNs. You might notice a long delay before the EFI Boot Manager menu screen is presented and when entering the EFI Shell, or before the device mappings are displayed and the EFI shell prompt is presented.

The following example shows the policy configuration dialog. In this example, the policy is being unchanged from the default policy.

```
# drvcfg -s
HP AVIO Stor Driver Configuration
=================================

Warning: enumerating all SCSI or FC LUNs increases initialization times.
```

```
Enumerate all SCSI LUNs (Y/N)? [current setting: N]: N

Enumerate all FC LUNs (Y/N)? [current setting: N

   Drv[2F] Ctrl[ALL] Lang[eng] - Options set. Action Required is None
None
None

Shell>

Reset the guest for the change to take effect
    vMP MAIN MENU
        CO: Console
        CM: Command Menu
        CL: Console Log
        SL: Show Event Logs
        VM: Virtual Machine Menu
        HE: Main Help Menu
        X: Exit Connection


[g1] vMP> CM

         (Use Ctrl-B to return to vMP main menu.)

[g1] vMP:CM>  RS
At next boot only boot LUN will be enumerated

    Use ^ and v to change option(s). Use Enter to select an option
Loading.: EFI Shell [Built-in]
EFI Shell version 1.10 [14.62]onsole - - - - - - - - - - -
Device mapping table
    fs0 : Acpi(PNP0A03,0)/Pci(0|0)/Scsi(Pun0,Lun0)/HD(Part1,SigBEC59C34-E6C8-11DB-8002-D6217B60E588)
    fs1 : Acpi(PNP0A03,0)/Pci(0|0)/Scsi(Pun0,Lun0)/HD(Part3,SigBEC59C70-E6C8-11DB-8004-D6217B60E588)
    blk0 : Acpi(PNP0A03,0)/Pci(0|0)/Scsi(Pun0,Lun0)
    blk1 : Acpi(PNP0A03,0)/Pci(0|0)/Scsi(Pun0,Lun0)/HD(Part1,SigBEC59C34-E6C8-11DB-8002-D6217B60E588)
    blk2 : Acpi(PNP0A03,0)/Pci(0|0)/Scsi(Pun0,Lun0)/HD(Part2,SigBEC59C52-E6C8-11DB-8003-D6217B60E588)
    blk3 : Acpi(PNP0A03,0)/Pci(0|0)/Scsi(Pun0,Lun0)/HD(Part3,SigBEC59C70-E6C8-11DB-8004-D6217B60E588)
startup.nsh> echo -off

    setting hpux path(\EFI\HPUX)...
    type 'fs[x]:' where x is your bootdisk (0, 1, 2...)
    type 'hpux' to start hpux bootloader

==============================================================================
```

For information about how to change the enumeration policy, see the *HP-UX vPars and Integrity VM V6.1 Administrator Guide*.

Changing the policy to enumerate all AVIO storage devices might result in longer guest boot times (to the EFI level), depending on the guest's configuration.

If you need to boot from a tape device attached to an NPIV (such as performing tape-based Ignite-UX recovery), change the enumeration policy to "Enumerate all FC LUNs". As mentioned previously, enumerating all FC LUNs can result in a long guest boot time. To minimize this delay, you can temporarily reduce the number of NPIV HBAs for the VM to only the one containing the tape boot device.

3.1.21 Probe of NPIV HBAs for Fibre Channel Targets Might Timeout

During online migration, the target VSP probes any NPIV HBAs for Fibre Channel (FC) targets. If this probe takes close to 30 seconds (or longer), online migration times out and aborts. The `hpvmmigrate` command prints an error:

```
# hpvmmigrate -P guest -h host2 -o
...
hpvmmigrate: Frozen phase (step 23) - progress 21%
    Target: (protocol low) header read timeout, 30 seconds
    Target: could not receive message header
```

On the target VSP machine, the syslog will contain the following warning message:

```
vmunix: HVSD: HPVM online migration warning: NPIV probe took too long, 79 seconds
```

To solve this problem, use the following workaround:

1. Verify that the probe time of each individual FC HBA is under 10 seconds.
2. On the target VSP, run the `ioscan` command to measure how long each FC HBA takes to probe. The ms_scan_time column gives the probe time.

```
# ioscan -P ms_scan_time -C fc
Class     I  H/W Path        ms_scan_time
=========================================
fc        0  44/0/0/0/0/0/0  0 min 7 sec 7 ms
fc        1  44/0/0/2/0/0/0  0 min 6 sec 213 ms
fc        2  44/0/0/2/0/0/1  0 min 4 sec 1 ms
```

3. If an individual FC HBA takes more than 10 seconds to probe, check your FC switch and zone settings to see why the probe time is so high.
4. Increase the online migration timeout value using the `hpvmmodify` command. For example, if the `syslog` reports a warning that the NPIV probe time is 79 seconds, increase the timeout value several seconds beyond that, to around 90 seconds (90000 msec):

```
# hpvmmodify -P guest1 -x "tunables=ogmo=90000"
# hpvmmodify -P guest1 -x migrate_frozen_phase_timeout=90
```

The default timeout value for ogmo is 30000 msec (30 seconds). The default timeout value for migrate_frozen_phase_timeout is 60 seconds.

Alternatively, you can also reduce the NPIV probe time by reducing the number of NPIV HBAs assigned to the guest.

3.1.22 vPar might hang during boot

A vPar might hang during boot, just prior to running the system run level startup `rc` scripts. This is expected to be a rare occurrence. Pressing the **Enter** key or any other key on the vPar (guest) console should resume boot up of the vPar. The problem does not occur with virtual machines.

3.1.23 Memory allocation might not be as expected

vPar and VM memory is allocated in granules of 64 MB. If a vPar's or VM's memory size is not an even multiple of 64 MB, the API/CLI code rounds up, but hpvmapp rounds down. Consequently, there will be 64 MB of memory reserved for the vPar/VM, but not used by it, and the vPar/VM might have up to 64MB to 1KB of memory less than was allocated.

To work around this problem, set the vPar or VM memory size to a multiple of 64 MB.

3.1.24 I/Os take long to complete under heavy I/O conditions on vPars/VMs with large NPIV LUN configuration

In a large LUN configuration, spread NPIV HBAs across multiple physical HBA ports at the VSP level.

3.1.25 vPar/VM hangs on boot just prior to RC scripts being launched

The vPar/VM might hang on boot just prior to RC scripts being launched, The boot hang is caused by a vPar/VM console output buffer overflow condition.

In the event of a vPar or VM hangs during boot, press the **Enter** key or any other key to resume the vPar/VM boot process.

3.1.26 CPU deconfigured by firmware can be assigned to vPar/VM

If a CPU experiences local MCA's too many times, the system firmware de-configures the CPU without HP-UX knowledge. This might cause the VSP to give the CPU to one of the vPars or VMs (which originally owned the CPU). The result of this is a vPar/VM hang in the bootloader.

The likelihood of this situation occurring is low, but if it does, stop and restart the vPar/VM to cause different CPUs to be assigned to the vPar or VM.

3.1.27 Probe of NPIV HBAs for Fibre Channel Targets May Timeout

During online migration, the target VM Host probes any NPIV HBAs for Fibre Channel (FC) targets. If this probe takes close to 30 seconds (or longer), online migration times out and aborts. The `hpvmmigrate` command prints an error:

```
# hpvmmigrate -P guest -h host2 -o
...
hpvmmigrate: Frozen phase (step 23) - progress 21%
    Target: (protocol low) header read timeout, 30 seconds
    Target: could not receive message header
```

On the target VM Host machine, the syslog will contain the following warning message:

```
vmunix: HVSD: HPVM online migration warning: NPIV probe took too long, 79 seconds
```

To solve this problem, use the following workaround:

1. Verify that the probe time of each individual FC HBA is under 10 seconds.
2. On the target VM Host, run the `ioscan` command to measure how long each FC HBA takes to probe. The ms_scan_time column gives the probe time.

   ```
   # ioscan -P ms_scan_time -C fc
   Class     I  H/W Path        ms_scan_time
   ========================================
   fc        0  44/0/0/0/0/0/0  0 min 7 sec 7 ms
   fc        1  44/0/0/2/0/0/0  0 min 6 sec 213 ms
   fc        2  44/0/0/2/0/0/1  0 min 4 sec 1 ms
   ```

3. If an individual FC HBA takes more than 10 seconds to probe, check your FC switch and zone settings to see why the probe time is so high.
4. Increase the online migration timeout value using the `hpvmmodify` command. For example, if the syslog reports a warning that the NPIV probe time is 79 seconds, increase the timeout value several seconds beyond that, to around 90 seconds (90000 msec):

   ```
   # hpvmmodify -P guest1 -x "tunables=ogmo=90000"
   # hpvmmodify -P guest1 -x migrate_frozen_phase_timeout=90
   ```

The default timeout value for ogmo is 30000 msec (30 seconds). The default timeout value for migrate_frozen_phase_timeout is 60 seconds.

Alternatively, you can also reduce the NPIV probe time by reducing the number of NPIV HBAs assigned to the guest.

3.1.28 vPars known behaviors and workarounds

The following table lists known vPars behaviors and workarounds.

Table 2 Known behaviors in vPars V6.1

Problem	Description/Action
When the requirements for satisfying Locality-Optimized Resource Alignment requirements for satisfying LORA configuration are not met, configuration are not met, a message is displayed.	Resources for a given vPar are chosen (as part of the vPar's boot up sequence) based on the resource availability and locality considerations. When the requirements for satisfying LORA configuration are not met, the following message is displayed in the syslog of the vPar OS instance: `AutoSensor: LORA Mode Disabled -- Single Locality Partition.` This message is seen even on a multiple locality vPar when it does not meet the LORA requirements. **What to do** You can ignore the message. For more details, see the Locality-Optimized Resource Alignment white paper on the BSC website http://h20000.www2.hp.com/bc/docs/support/SupportManual/c02070810/c02070810.pdf
No vPar crash dump when the VSP crashes on TOC/MCA	When a TOC is performed on the VSP or the VSP encounters an MCA, no vPar crash dump is provided. Only a VSP crash dump is available and it is augmented with some debugging information from the vPars. **What to do** Currently, there is no workaround.
CPU/memory info in `machinfo` output on VSP could be confusing	The `machinfo` command displays system information from the HP-UX view of the system configuration. The `machinfo` command might show different values based on when the command is executed. If executed on the VSP after installing HP-UX and before installing the Integrity VM software, `machinfo` shows all the sockets and logical processors. The logical processor count represents cores if the kernel tunable `lcpu_attr` value is 0 and threads when `lcpu_attr` value is 1. You can obtain the value of `lcpu_attr` by using the `kctune` command. Note that `lcpu_attr` is set to zero in the VSP by default for optimal VSP performance, and so, the logical processor count is always the CPU core count. After the Integrity VM software is installed, the logical processor count of `machinfo` represents the number of VSP logical processors and the logical processors in the vPar/VM pool, but not yet activated in a vPar. When a vPar is started, the logical processors in the vPar/VM pool assigned to the vPar are deallocated from the VSP and the `machinfo` output in the VSP will reflect that reduction in logical processor count. When a vPar is stopped, the processor count shown in the VSP `machinfo` output will increase by the number of CPUs assigned to the vPar. The memory value displayed by the `machinfo` command shows the amount of memory that was available to HP-UX when booted on the VSP. This memory value includes memory that is allocated to the vPars and the memory used by the VSP. Unlike the logical processor count, the memory amount does not change with the installation of the Integrity VM software. **What to do** As workaround, use the `vparhwmgmt –p cpu –l` command to view the number of processor cores that are allocated to the VSP and to the vPar pool.

Table 2 Known behaviors in vPars V6.1 *(continued)*

Problem	Description/Action
vm.core will not be generated when vPar is in OS upon TC	A vPar that transitions past the load of the HP-UX kernel will not produce a vm.core on the VSP when it is terminated abnormally via a TOC. Once HP-UX vm.core will not be generated when vPar is in OS upon TC transitions to the boot state, the responsibility for state capture when a TOC occurs is via DUMP. An activating vPar that has not transitioned to the boot state produces a vm.core when a TOC occurs. **What to do** There is no workaround. This is an expected behavior in vPars.
CPU OLD Policy	The vPars v6 product supports dynamic CPU addition and deletion. The selection criteria for CPU addition are performed from within the VSP based on LORA CPU OLD policy policies. The selection criteria for CPU deletion is performed from within the HP-UX instance that is target of the CPU deletion. The following criteria are used to select the CPU cores to be deleted: • Only cores in the default pset (see `psrset(1M)`) can be dynamically removed. • The monarch core can never be deleted. • If the default pset does not contain enough cores to satisfy a full request to reduce the number of cores, the processor assignments will remain unchanged. • Processors can be moved to the default pset and then deleted. **What to do** There is no workaround.
TC/RS/stop behavior during Serviceguard TOC	When Serviceguard detects a hang of the operating system, it uses a system interface to TOC the operating system. If Serviceguard initiates a TOC in a TC/RS/stop behavior during Serviceguard TOC v6.0 vPar, any subsequent `vparreset` operation initiated from the VSP to that specific vPar causes a message to be logged to `syslog`. The command completes but no action is taken until the Serviceguard TOC operation completes. **What to do** There is no workaround.
NPIV LUNs not shown by default invocation of ioscan	An ioscan issued from within a vPar does not display any LUNs behind the NPIV HBA unless the –N option is specified. The ioscan command without NPIV LUNs not shown by default invocation of ioscan the –N option only displays devices that use the old style device file format. The NPIV LUNs use the agile device file format and require the `–N` option to `ioscan` in order to display LUNs in the output. **What to do** Use the `–N` option with the `ioscan` command to view devices behind NPIV HBAs.

Table 2 Known behaviors in vPars V6.1 *(continued)*

Problem	Description/Action
Cannot remove a VLAN-based vNIC if the VLAN has been removed	If a vPar was created with a vswitch based on a VLAN, and that VLAN has been removed, it is not possible to modify or remove a vNIC connected to that vswitch. **What to do** Recreate the VLAN with `nwmgr` and restart the vswitch. This will allow you to remove the vNIC using the `vparmodify` command.

Because the base page size of VSP is set to 64K, the `vmstat` command might show incorrect results	If the base page size is changed to value other than 4K, the output value corresponding to 4K is the value for the new base page size. This is true for further displayed page sizes i.e. 8K, 16K, 64K, and so on.
	Example 1 When base page size = 4K
	`...`
	`149003 Page Select Size Successes for Page size 4K`
	`139682 Page Select Size Successes for Page size 8K`
	`494477 Page Select Size Successes for Page size 16K ...`
	Example 2 When base page size = 64K
	`...`
	`28 Page Select Size Successes for Page size 4K`
	`67 Page Select Size Successes for Page size 16K`
	`14 Page Select Size Successes for Page size 64K`
	`2 Page Select Size Successes for Page size 256K ...`
	In the example for base page size 64, the value 28 actually refers to the page size 64K, 67 refers to the page size 256K, and so on.
	What to do
	There is no workaround.

Table 2 Known behaviors in vPars V6.1 *(continued)*

Problem	Description/Action

Because the base page size of host is set to 64K, the following commands might show incorrect results:

- acctcom
- acctcms
- acctprc2

These commands show incorrect values for mean memory size and kcore-minutes when the base page size is altered to any value other than 4K. In order to see the correct value, the value should be multiplied by (base page size / 4K).

Example 3 Base page size= 64K

```
# /usr/sbin/acct/acctcom -mk mypacct3
COMMAND START END REAL CPU MEAN KCORE
NAME USER TTYNAME TIME TIME (SECS) (SECS) SIZE(K) MIN
#accton root pts/ta 15:42:33 15:42:33 0.01 0.01 0.00 0.00
dd root pts/ta 15:42:33 15:42:33 0.05 0.04 120.00 0.08
dd root pts/ta 15:42:33 15:42:33 0.06 0.04 120.00 0.08
dd root pts/ta 15:42:33 15:42:33 0.07 0.08 120.00 0.16
.
.
.
```

The actual value for KCORE-MIN will become (0.08 * (64/4)). The actual value for MEAN SIZE will become (120 * (64/4)).

```
# /usr/sbin/acct/acctcms -a mypacct3
TOTAL COMMAND SUMMARY
COMMAND NUMBER TOTAL TOTAL TOTAL MEAN MEAN HOG CHARS BLOCKS
NAME CMDS KCOREMIN CPU-MIN REAL-MIN SIZE-K CPU-MIN FACTOR TRNSFD READ
TOTALS 11 3.06 0.02 0.03 127.53 0.00 0.87 4290776228 0
dd 10 3.06 0.02 0.03 127.53 0.00 0.88 4290772992 0
accton 1 0.00 0.00 0.00 0.00 1.00 6000.00 3236 0
```

The actual value for TOTAL KCOREMIN will become (3.06 * (64/4)). The actual value for MEAN SIZE-K will become (127.53 * (64/4)).

What to do

There is no workaround.

Table 2 Known behaviors in vPars V6.1 *(continued)*

Problem	Description/Action
Interrupt balancing daemon is disabled in vPars V6.1	Frequent interrupt migration on the vPar can lead to storage LUNs going offline. This can occur indirectly when dynamic CPU migration occurs on a frequent basis when the interrupt balancing daemon is enabled. **What to do** The interrupt balancing daemon is disabled by default. Do not enable the interrupt balancing daemon.

Not enough memory allocated for VSP when the VSP base page size is set to 4K.	If the VSP base page size is changed to 4K from the default 64K, the VSP might run into low memory situation. Run `hpvmstatus -s` to check the current VSP free memory. For example:
	`# hpvmstatus -s`
	`…`
	`Available VSP memory = 2139 Mbytes`
	What to do
	Increase the memory allocated for VSP by tuning the
	`HPVM_MEMORY_OVERHEAD_PERCENT`
	variable in `/etc/rc.config.d/hpvmconf` and then reboot the VSP. For example, run the following commands to reserve around 15% of system free memory as VSP memory:
	`# ch_rc -a -p HPVM_MEMORY_OVERHEAD_PERCENT=15.0 /etc/rc.config.d/hpvmconf`
	`# shutdown -r now`

4 Integrity VM Support Policy

This chapter describes the HP-UX vPars and Integrity VM support policies and software version requirements for Integrity VSP and guest operating system environments.

4.1 vPars and Integrity VM Minimum Support Life

The following table shows the support date for vPars and Integrity VM versions.

Integrity VM Version	**Release Date**	**Expected End of Support Date**	**Current Status**
A.03.50	December 2007	Same as HP-UX 11i v2	Supported
B.04.10	April 2009	April 2012	Supported
B.04.20	March 2010	March 2013	Supported
B.04.20.05	September 2010	September 2013	Supported
B.04.30	March 2011	December 2013	Supported
vPar Version	**Release Date**	**Expected End of Support Date**	**Current Status**
B.06.00	November 2011	November 2014	Supported
vPars and Integrity VM Version	**Release Date**	**Expected End of Support Date**	**Current Status**
B.06.10	March 2012	December 2014	Supported

4.2 vPars Firmware Requirements

- Requirement:

  Customers must upgrade to the latest release in order to receive defect fixes.

  vPar system firmware requirements: The latest available firmware for the BL8x0c i2 servers that is available can be installed. Firmware bundle 2011.05 is the minimum requirement. For Superdome 2 servers, Superdome 2 Dynamic Core Firmware (2.51.102) is required.

NOTE: Customers using only VMs have no firmware requirements.

4.3 VSP OS and Server Support

	Integrity VSP or Server Support	**Integrity VM Version A.03.50**	**Integrity VM Version B.04.10**	**Integrity VM Version B.04.20**	**Integrity VM Version B.04.20.05**	**Integrity VM Version B.04.30**	**vPars and Integrity VM B.06.10**	**Notes**
	VSP OS Support	HP-UX 11i v2 December 2007	HP-UX 11i v3 March through September 2009	HP-UX 11i v3 March 2010	HP-UX 11i v3 September 2010	HP-UX 11i v3 March 2011 through September 2011	HP-UX 11i v3 March 2012	New vPars and Integrity VM versions support the latest HP-UX OE release at the time of the vPars and Integrity VM release. For more information about support for specific HP-UX versions, see an authorized HP representative.
VSP Server Support	The vPars and Integrity VM product supports all Integrity servers. New servers are supported on the latest Integrity VM version concurrent with the shipment of that server.							

For information about installing vPars and Integrity VM V6.1, see the *HP-UX vPars and Integrity VM V6.1 Administrator Guide*.

4.4 HP-UX vPar and VM Support

Type of Guest OS Support	Integrity VM Version A.03.50	Integrity VM Version B.04.10	Integrity VM B.04.20	Integrity VM B.04.20.05	Integrity VM B.04.30	vPars and Integrity VM B.06.10	Notes
HP Integrity Servers	HP-UX 11i v3 March 2007 through September 2008	HP-UX 11i v3 March 2007 through September 2009	HP-UX 11i v3 September 2007 through March 2010	HP-UX 11i v3 March 2008 through September 2010	HP-UX 11i v3 March 2010 through September 2011	Integrity VM — HP-UX 11i v3 March 2010 through March 2012 vPars — Not supported	
HP Integrity server blades	Not supported	HP-UX 11i v3 March 2007 through September 2009	HP-UX 11i v3 September 2007 through March 2010	HP-UX 11i v3 March 2008 through September 2010	HP-UX 11i v3 March 2010 through September 2011	Integrity VM — HP-UX 11i v3 March 2010 through March 2012 vPars — Not supported	
HP Integrity server blades (i2–based)[1]	Not supported	Not supported	HP-UX 11i v3 March 2009 through March 2010[2]	HP-UX 11i v3 March 2009 through September 2010	HP-UX 11i v3 March 2010 through September 2011	Integrity VM — HP-UX 11i v3 March 2010 through March 2012 vPars — HP-UX 11i v3 September 2011 through March 2012	
HP Integrity Superdome 2[1]	Not supported	Not supported	Not supported	HP-UX 11i v3 September 2010	HP-UX 11i v3 March 2010 through September 2011	Integrity VM — HP-UX 11i v3 September 2010 through March 2012 vPars — HP-UX 11i v3 March 2012	
HP Integrity rx2800 i2 Servers	Not supported	Not supported	Not supported	HP-UX 11i v3 September 2010	HP-UX 11i v3 March 2010 through September 2011	Integrity VM — HP-UX 11i v3 September 2010 through March 2012 vPars — HP-UX 11i v3 March 2012	

[1] Support for the latest Intel® Itanium® Processor 9300 series

[2] Requires V4.2 with the following VM Host patches:

PHSS_40875 1.0 HPVM B.04.20 CORE PATCH
PHSS_40876 1.0 HPVM B.04.20 VMAGENT
PHSS_40901 1.0 HPVM B.04.20 VMMIGRATE PATCH

NOTE: For information about required patches, see Chapter 2 (page 11).

HP-UX patches might be required for proper operation. Install these patches on HP-UX guest operating systems as necessary. For specific patch information, see the patches listed in Chapter 2 (page 11) .

4.5 Storage Interface Support

vPars and Integrity VM V6.1	Notes
• Fibre Channel adapters supported by TD, FCD, or FLCP drivers • SCSI adapters supported by the C8xx, MPT, or CISS drivers • IDE adapters supported by the SIDE driver • USB support of the UsbScsiAdapter driver and USB 2.0 support of the usb_ms_scsi driver (AVIO HVSD support for USB 2.0 as of 11i v3 0903) • iSCSI adapters supported by the iSCSI driver (AVIO requires HostAVIOStor B.11.31.1009.01 or later) • SAS adapters supported by the SASD driver • For B.06.10 and AVIO only, FCoE support provided by CNA supported by the FCoC driver • SATA DVD is supported by AVIO only. Even if SATA DVD is set as a SCSI device, the guest EFI does not recognize it.	Accelerated virtual I/O (AVIO) storage interfaces are defined using the *avio_stor* adapter type. (See the `hpvmresources` manpage.)

4.5.1 vPar and VM Attached Device Support

vPars and Integrity VM Version V6.1	Notes
• CD/DVD burners • Media changers • Tape devices	Attached devices are supported for all types of guest operating systems that provide supported drivers for the physical device as attached to the VSP device.

4.5.2 Multipathing Software Support

Multipathing software is supported on the Integrity VSP system. Unless specified otherwise, multipathing is supported for use with AVIO (virtual adapter type specified as *avio_stor*). The required version for these products is determined by the software vendor and the release of HP-UX installed on the VSP system.

Backing Store	vPars and Integrity VM Version 6.1
Whole Disk (or LUN)	• HP-UX 11i v3 built-in multipathing • EMC PowerPath[1] with legacy DSFs
LVM Logical Volumes	• HP-UX 11i v3 built-in multipathing • PVLinks • EMC PowerPath[1] with legacy whole disk DSF in a Volume Group
VxVM Logical Volumes	• HP-UX 11i v3 built-in multipathing • EMC PowerPath[1] with legacy whole disk DSF in a Disk Group • Symantec DMP
VxFS Files	• HP-UX 11i v3 built-in multipathing • PVLinks • EMC PowerPath[1] with legacy disk DSF in a Volume Group • Symantec DMP

[1] Supported by EMC. See EMC documentation.

4.5.3 EVA Series Firmware Requirement

The AVIO Storage Driver supports Active-Active firmware types on EVA series (3000/5000 and 4000/6000/8000). Be sure that the following firmware revision levels are met before configuring AVIO backing stores on these arrays:

- On EVA 4000/6000/8000, all released firmware revisions support Active-Active configuration. Therefore, no action is necessary on these arrays.
- EVA 3000/5000 arrays need minimum firmware revisions of VCS v4.004 (Active-Active firmware).

For more information about EVA firmware upgrades, see the HP Services website: HP Services.

4.6 NPIV configuration support

Limit Description	Supported Limit
NPIV HBAs per vPar/VM	8
Number of NPIV HBAs per physical HBA	16
Number of paths supported per NPIV device	8
Number of LUNs per NPIV HBA	2048
Number of NPIV devices per vPar/VM	2048

NOTE: In configurations where multiple NPIV HBAs created from a single physical HBA are used by different vPars/VMs, all the I/O from these vPars/VMs share a single physical HBA, which could lead to performance bottlenecks in high I/O scenarios.

For a more balanced performance, HP recommends that you spread NPIV HBAs for vPars/VMs across multiple physical adapters.

4.7 Network interface support for AVIO

Specific network interfaces are supported if they are supported for the version of HP-UX in use on the VSP system. The VSP physical network interface card can be configured with Auto Port Aggregation (APA), with the resulting port supported as indicated in the following table:

Virtual Network Adapter Type	vPars and Integrity VM V6.1
avio_lan (AVIO)	The following Ethernet drivers are supported, including APA ports: • iether • igelan • ixgbe • icxgbe • iexgbe • Networking support provided by CNA, supported by the iocxgbe driver

4.8 Supported adapters for AVIO networking

vPars and Integrity VM supports those adapters that are of Ethernet or the IEEE 802.3 CSMA/CD network. Note that AVIO interfaces are supported by a select set of Ethernet host NICs. The following Ethernet cards are supported with AVIO on HP-UX vPars and VMs:

- A6794A — Core
- A6825A — Core
- A7109A — Core
- A6847A — Add-in
- AB465A — Combo Add-in
- A9782A — Combo Add-in
- A9784A — Combo Add-in
- AB352 — Core
- AB545A — Networking Add-in
- A7011A — Networking Add-in
- A7012A — Networking Add-in
- AB290A — Combo Add-in
- AB287A — Add-in (10GbE)
- AD331A — Add-in
- AD332A — Add-in
- AD193A — Combo Add-in
- AD194A — Combo Add-in
- AD221A — Add-in
- AD222A — Add-in
- AD337A — Add-in
- AD338A — Add-in
- AD339A — Add-in
- AD385A — Add-in (10GbE)
- AD386A — Add-in (10GbE)
- HBA's AM225A, AM232A, AM233A
- NC360m — C-class Mezz card
- NC364m — C-class Mezz card
- HP NC551m Dual Port FlexFabric 10Gb Converged Network Adapter

4.9 Network interface support for direct I/O networking

The direct I/O networking feature has the following hardware, firmware, and driver support requirements:

Platforms:

- HP Integrity rx2800 i2
- HP Integrity BL8x0c i2
- HP Integrity Superdome 2

Adapters:

- Broadcom: Device Level Assignment (DLA)
  - LOMs: BL8x0c i2, Superdome 2
  - Stand-Up Cards: AM225A/AM232A/AM233A
  - Mezzanine Cards: NC532m
  - Broadcom cards must be running this firmware version or later: ROM Firmware (Boot code) Version: 6.2.21

    NOTE: You can obtain the ROM firmware from HP.COM or the IPF offline CD 1203.

  - Broadcom VSP and vPars/VMs must be running this minimum driver version: 10GigEthr-02 (iexgbe) B.11.31.1203.
- Emulex: Function Level Assignment (FLA)
  - Stand-Up Cards: AT118A
  - CNA Stand Up (NIC only): AT111A
  - Mezzanine Cards: NC552m
  - CNA Mezzanine (NIC only): NC553m
  - Emulex cards must be running this firmware version or later: ROM Firmware Version : 3.X04.356.3
  - Emulex VSP and vPars/VMs must be running this minimum driver version: 10GigEthr-03 (iocxgbe) B.11.31.1203

NOTE: All drivers must be from HP-UX 11i v3 March 2012 or later.

4.10 Online and Offline Migration Support

The following list provides the supported online and offline migration paths for HP-UX guests:

- Support:
  - Forward migration — Supported between two VSPs running the same or subsequent vPars and Integrity VM version (see Table 3 (page 34)).
  - Backward migration — Supported for a vPar/VM migrating back only to a VSP running the same version on which the vPar/VM was originally booted (including any vPar and Integrity VM release patches that were installed on the original VSP) and has not been rebooted since. As of Integrity VM Version 4.3, a VM can also be migrated back to a VSP that is running a later version of Integrity VM on which the VM was originally booted. For example, a VM booted on a VSP running Integrity VM Version 4.2.5, then migrated

to a VSP running Version 4.3, can be migrated back to a VSP running version 4.2.5 (provided the VM has not been rebooted since beginning its migration).

◦ **NOTE:** In addition to the restrictions shown in the following online migration path table, the processors on the source and destination VSPs must all be within one of the following groups:

– All Integrity VM supported variants of the Itanium 2 processor prior to the 9000
– Itanium 2 9000 and the Itanium 2 9100
– Itanium 9300

Online forward migration table:

Table 3 Integrity VM Online Forward Migration Paths

Integrity VM Version	Supported Forward Migration Path
Integrity VM V4.1	Integrity VM V4.1 or Integrity VM V4.2
Integrity VM V4.2	Integrity VM V4.2, Integrity VM V4.2.5, or Integrity VM4.3
Integrity VM V4.2.5	Integrity VM V4.2.5 or Integrity VM V4.3
Integrity VM V4.3	Integrity VM V4.3 or Integrity VM V6.1

- Offline migration support:
  ◦ Offline migration support table:

Table 4 Offline Forward and Backward Migration Support Paths

vPars and Integrity VM Version	Supported Offline Migration Path (Forward and Backward)
Integrity VM V3.5	Integrity VM V3.5
Integrity VM V4.1 or later	Integrity VM V4.1 or later
vPars V6.0	vPars V6.1

NOTE: When migrating across versions (Integrity VM 4.1 and later), be aware of the vPar/VM OS versions that are supported on the target VSP. If the vPar/VM is not supported by the target VSP, the vPar/VM must be updated immediately after migrating to maintain support on the target VSP.

5 Support and other resources

5.1 Contacting HP

5.1.1 Before you contact HP

Be sure to have the following information available before you call contact HP:

- Technical support registration number (if applicable)
- Product serial number
- Product model name and number
- Product identification number
- Applicable error message
- Add-on boards or hardware
- Third-party hardware or software
- Operating system type and revision level

5.1.2 HP contact information

For the name of the nearest HP authorized reseller:

- See the Contact HP worldwide (in English) webpage (http://welcome.hp.com/country/us/en/wwcontact.html).

For HP technical support:

- In the United States, for contact options see the Contact HP United States webpage (http://welcome.hp.com/country/us/en/contact_us.html). To contact HP by phone:
  - Call 1-800-HP-INVENT (1-800-474-6836). This service is available 24 hours a day, 7 days a week. For continuous quality improvement, calls may be recorded or monitored.
  - If you have purchased a Care Pack (service upgrade), call 1-800-633-3600. For more information about Care Packs, see the HP Care Packs website at HP Care Packs.
  - In other locations, see the Contact HP worldwide (in English) webpage (http://welcome.hp.com/country/us/en/wwcontact.html)

5.1.3 HP Insight remote support

HP strongly recommends that you install HP Insight Remote Support software to complete the installation or upgrade of your product and to enable enhanced delivery of your HP Warranty, HP Care Pack Service or HP contractual support agreement. HP Insight Remote Support supplements your monitoring, 24x7 to ensure maximum system availability by providing intelligent event diagnosis, and automatic, secure submission of hardware event notifications to HP, which will initiate a fast and accurate resolution, based on your product's service level. Notifications may be sent to your authorized HP Channel Partner for on-site service, if configured and available in your country. The software is available in two variants:

- HP Insight Remote Support Standard: This software supports server and storage devices and is optimized for environments with 1-50 servers. Ideal for customers who can benefit from proactive notification, but do not need proactive service delivery and integration with a management platform.
- HP Insight Remote Support Advanced: This software provides comprehensive remote monitoring and proactive service support for nearly all HP servers, storage, network, and SAN environments, plus selected non-HP servers that have a support obligation with HP. It is

integrated with HP Systems Insight Manager. A dedicated server is recommended to host both HP Systems Insight Manager and HP Insight Remote Support Advanced.

Details for both versions are available at:

http://www.hp.com/go/insightremotesupport

To download the software, go to Software Depot:

http://www.software.hp.com

Select Insight Remote Support from the menu on the right.

NOTE: HP recommends using Insight Remote Support on the VSP system. Information from Insight Remote Support running on virtual machines should not be used to determine the hardware state.

5.1.4 Subscription service

HP recommends that you register your product at the Subscriber's Choice for Business website: http://www.hp.com/country/us/en/contact_us.html. After registering, you will receive an email notification of product enhancements, new driver versions, firmware updates, and other product resources.

5.1.5 Documentation feedback

HP welcomes your feedback. To make comments and suggestions about product documentation, send a message to **docsfeedback@hp.com**. Include the document title and manufacturing part number. All submissions become the property of HP.

5.2 Related information

The following documents [and websites] provide related information:

Table 5 Documentation and its location

Documents	Website
HP-UX vPars and Integrity VM V6.1 Administrator Guide	http://www.hp.com/go/virtualization-manuals
HP-UX GUID Manager Administrator Guide	http://www.hp.com/go/hpux-vpars-docs and http://www.hp.com/go/insightdynamics-manuals
• *HP Integrity Virtual Server Manager User Guide* • *HP Integrity Virtual Server Manager Release Notes* **NOTE:** The HP Integrity Virtual Server Manager is the graphical user interface version of the command-line interface HP-UX vPars V6.1.	http://www.hp.com/go/matrixoe/docs and http://www.hp.com/go/insightdynamics-manuals
• *HP BladeSystem Onboard Administrator Command Line Interface User Guide Version 3.30* • *HP BladeSystem Onboard Administrator User Guide Version 3.30* • *HP BladeSystem c3000 Enclosure* (whitepaper) • *HP BladeSystem c7000 Enclosure technologies* (whitepaper)	http://www.hp.com/go/blades_enclosures-docs In the main page, click **HP BladeSystem c3000 Enclosures** or **HP BladeSystem c7000 Enclosures**.
Virtual Partitions documentation	www.hp.com/go/hpux-vpars-docs

Table 5 Documentation and its location *(continued)*

Documents	Website
HP-UX 11i v3 documentation	http://www.hp.com/go/hpux-core-docs In the main page, click **HP-UX 11i v3**.
HP Serviceguard Toolkit for Integrity Virtual Servers User Guide	HP Serviceguard Toolkit for Integrity Virtual Servers User Guide

5.3 Typographic conventions

This document uses the following typographical conventions:

`%`, `$`, or `#`	A percent sign represents the C shell system prompt. A dollar sign represents the system prompt for the Bourne, Korn, and POSIX shells. A number sign represents the superuser prompt.
audit(5)	A manpage. The manpage name is *audit*, and it is located in Section 5.
`Command`	A command name or qualified command phrase.
`Computer output`	Text displayed by the computer.
Ctrl+x	A key sequence. A sequence such as **Ctrl+x** indicates that you must hold down the key labeled **Ctrl** while you press another key or mouse button.
`ENVIRONMENT VARIABLE`	The name of an environment variable, for example, `PATH`.
ERROR NAME	The name of an error, usually returned in the `errno` variable.
Key	The name of a keyboard key. **Return** and **Enter** both refer to the same key.
Term	The defined use of an important word or phrase.
`User input`	Commands and other text that you type.
`Variable`	The name of a placeholder in a command, function, or other syntax display that you replace with an actual value.
[]	The contents are optional in syntax. If the contents are a list separated by `|`, you must choose one of the items.
{}	The contents are required in syntax. If the contents are a list separated by `|`, you must choose one of the items.
...	The preceding element can be repeated an arbitrary number of times.
□	Indicates the continuation of a code example.
|	Separates items in a list of choices.
WARNING	A warning calls attention to important information that if not understood or followed will result in personal injury or nonrecoverable system problems.
CAUTION	A caution calls attention to important information that if not understood or followed will result in data loss, data corruption, or damage to hardware or software.
IMPORTANT	This alert provides essential information to explain a concept or to complete a task
NOTE	A note contains additional information to emphasize or supplement important points of the main text.