

Top Ten Tips for using Virtual Partitions

A white paper from HP

Alan E. Hymes



Introduction	3
#1 – Kernel Memory Allocation (HP9000 only)	3
Symptoms: Performance, Functionality	3
Prevention and Suggestions	3
#2 – nPars and vPars Mode (Integrity only)	3
Symptoms: Functionality	3
Prevention and Suggestions	4
#3 – Memory Granularity and Booting	4
Symptoms: Functionality	4
Prevention and Suggestions	4
#4 – Memory Granularity and Performance	5
Symptoms: Performance	5
Prevention and Suggestions	5
#5 – Number of vPars and Cells per nPar	5
Symptoms: Supportability	5
Prevention and Suggestions	5
#6 – Mixing vPar and HP-UX release versions	6
Symptoms: Supportability	6
Prevention and Suggestions	6
#7 – vPars and I/O Interrupt Processing	6
Symptoms: Performance	6
Prevention and Suggestions	7
#8 – Memory Overhead with vPars	7

Symptoms: Performance, Functionality	7
Prevention and Suggestions	7
#9 – vPars and Securepath	7
Symptoms: Performance, Supportability	7
Prevention and Suggestions	7
#10 – vPars and OLAR.....	8
Symptoms: Functionality.....	8
Prevention and Suggestions.....	8

Introduction

HP-UX Virtual Partitions (vPars) has been available for about six years now. During this time most of the features and “rules” have not changed. However with more and more customers deploying vPars especially in a production environment there are implementation considerations. These are often overlooked when upgrading or expanding existing servers that are already running vPars. These considerations are mainly in three general areas:

- **Performance** – since vPars is para-virtualized there are no special considerations for the performance of a vPar. However, there are configuration issues that do impact the startup and shutdown of vPars.
- **Functionality** – functionality here refers more to configuration issues due to the fact that vPars does not share hardware below the Local Bus Adapter (LBA) level or PCI Slot level. Configuration of memory and/or granularity also effects functionality.
- **Supportability** – lastly, HP has set support “rules” which must be followed in order for the complex configuration to be supported. These “rules” have been determined by the vPar’s engineering team based on lab tests and code design.

This white paper outlines the ten most common issues in vPar configurations that can lead to problems in the above categories. For each issue, the most likely symptoms will be listed using these categories (**Performance, Functionality, Supportability**).

#1 – Kernel Memory Allocation (HP9000 only)

Symptoms: Performance, Functionality

On PA-RISC, HP-UX has a requirement that a kernel’s text and data must fit within the lower 2GB address space and must be contiguous. The reason is that some older, more complicated assembly routines within the kernel perform address manipulation using 32-bit arithmetic. There are no plans to change this on PA.

This requirement also applies in a Virtual Partition’s environment with the addition that ALL vPars kernels’ text and data (added together) must fit under the 2GB lower address space and each kernel must reside in contiguous memory. This restriction usually isn't a problem unless you significantly increase some of the tunable parameters that affect the data segment. As the vPar monitor is loading kernels (booting vPars), if one or more of the kernels won’t fit, the load will fail with a message like this in the monitor log:

```
INFO:CPU2:MON:No error free contiguous memory in the system to load the kernel below 2GB
```

Prevention and Suggestions

There are several possible solutions which are documented in the HP Whitepaper titled “How kernel memory is allocated and may be controlled in a vPars environment”. The best solution is to create each vPar defining the base and range of memory to be some value that you are sure the kernel will not grow beyond and make sure the total sum of all the memory allocated to the vPars is less than the total of the nPar. Please see the whitepaper for the details and example calculations.

#2 – nPars and vPars Mode (Integrity only)

Symptoms: Functionality

On an Integrity system, you must set the mode in order to boot into a specific mode. For

vPar's usage, there are only two modes:

- vPars - sets the next nPartition boot to boot into the vPars environment. This allows you to boot the vPars Monitor and therefore the virtual partitions in the next nPartition boot. You still need to boot the vPars Monitor and the virtual partitions, but this mode allows you to do this.
- nPars - sets the next nPartition boot to boot into the standalone environment. In this mode, you cannot boot the vPars Monitor and therefore the virtual partitions. However, you can boot any OE instance into standalone mode.

Prevention and Suggestions

You can set the mode from the following levels using the corresponding commands:

HP-UX Shell: vparenv [-m mode]

```
#vparenv -m vPars
```

Monitor: reboot [mode]

```
MON> reboot nPars
```

EFI: vparconfig [reboot mode]

```
fs0:\> vparconfig reboot vPars
```

Please see the document "Installing and Managing HP-UX Virtual Partitions" for more details.

#3 – Memory Granularity and Booting

Symptoms: Functionality

Granularity refers to the unit size in which memory is assigned to all virtual partitions in a given vPars database (vpdb). Granularity reflects only the unit size of memory and not the amount of memory that is assigned.

The default granularity is 128 MB for ILM and 128 MB for CLM. However, you can specify your own granularity for CLM and/or ILM. Granularity has some specific restrictions and cannot be changed in a vPars database after they are set.

Prevention and Suggestions

Granularity in the vPars database can only be specified during the creation of the vPars database. This means the first vparcreate command performed to create the database can be used to specify the granularity, but it cannot be changed after that. It cannot be changed by subsequent vparcreate commands nor any other commands; any change in values requires the entire vPars database to be re-created.

To change the granule sizes, it is necessary to:

1. record all virtual partition configurations
2. boot standalone
3. remove /stand/vpdb
4. re-create all virtual partitions, specifying the new granule sizes to the first vparcreate command

Please see the document "Installing and Managing HP-UX Virtual Partitions" for more details.

#4 – Memory Granularity and Performance

Symptoms: Performance

Coming up with a suitable granule size

Prevention and Suggestions

If your system (or nPartition) contains a large amount of memory (32 GB or more), you should consider the following when setting the granularity size:

- (a) VM performance: Virtual Memory subsystem in the OS cannot create pages bigger than the granule size. Run-time performance will be better on systems with higher page size.
- (b) Hitting nPar granule limits: Firmware supports certain max number of granules per nPar. If the system has more memory, granule size should be increased to fit within the limit.
- (c) Memory migration flexibility: Smaller granule sizes provide more flexibility in moving the memory around. So, customers may want to strike a balance between the flexibility and performance needs.
- (d) Long time to boot: With smaller granule sizes, vPars were taking a long time to boot on IPF systems. This problem is addressed by the patch PHKL_35698 in 11iv2 and this is not a problem in 11iv3. PHKL_35698 is part of 1123.0706 FEATURE11i bundle. vPars release A.04.04 requires the installation of 1123.0706 FEATURE11i bundle.
- (e) PA-RISC systems: ILM granule size must be small enough for each vPar to have at least one granule below 2GB for its memory, plus one granule for the monitor. 1GB granules will be too large for ILM on PA-RISC systems, therefore, and 128MB or 256MB granules will be better.

Please see the document “Installing and Managing HP-UX Virtual Partitions” for more details.

#5 – Number of vPars and Cells per nPar

Symptoms: Supportability

The current limit regarding the number of vPars and the number of cells per nPar in a vPar’s environment is 8 for both. So, the current requirement is a maximum of 8 vPars per nPar where the nPar is a maximum of 8 cells in a vPars environment. The cells may cross SuperDome cabinets, but the maximum number is 8.

Prevention and Suggestions

Although the vPar software currently allows you to create vPars outside these conditions, you will be in an unsupported state and may run into problems.

Please see the document “Ordering and Configuration Guide for Virtual Partitions” for more details.

#6 – Mixing vPar and HP-UX release versions

Symptoms: Supportability

Each system (on supported non-cell based) and each nPar (on supported cell based) has only one vPar Monitor running which manages all of the vPars in that partition. And each vPar has its own isolated Operating System. Due to the architecture of vPars and specific operating system functionality, there is a supported vPar version to OS version relationship that must be followed. This is primarily identified by the vPar Monitor since it is what is managing the whole vPar environment in a given hard partition.

Currently as of the releases shown below, only one version of HP-UX is supported with each major release of vPars: 11.11 (A.03.xx), 11.23 (A.04.xx) and 11.31 (A.05.xx). However, with vPars A.05, you may mix vPar versions in the same nPar as shown below.

Prevention and Suggestions

Currently the following environments are supported:

- A.03.xx Release Stream
 - Supports 11.11 / 11iv1 ONLY
 - Supports PA ONLY
 - Maintained as a separate release stream
 - Can NOT mix with other release streams (ie. A.04) in same Hard Partition

- A.04.xx Release Stream
 - Supports 11.23 (0505 & above) / 11iv2 (0505 & above) ONLY
 - Supports PA and IPF
 - Maintained as a separate release stream
 - Can NOT mix with other release streams (ie. A.03) in same Hard Partition

- A.05.xx Release Stream
 - A.05.xx monitor supports 11.31 (A.05.xx) and 11.23 (A.04.02)
 - Supports PA and IPF
 - Maintained as a separate release stream
 - Can mix with other release streams starting with A.04.02 in same Hard Partition

Please see the whitepapers called “Architectural Considerations when Mixing Virtual Partition and Operating System Versions” and “Mixed OS (HP-UX 11i v2/v3) Considerations in vPars A.05.xx” for more details.

#7 – vPars and I/O Interrupt Processing

Symptoms: Performance

vPars releases prior to A.04 have CPUs that can be configured as bound or unbound (floating) CPUs. The basic differences between a bound CPU and unbound CPU are that an unbound CPU does not participate in I/O interrupt processing.

In A.04 all CPUs are capable of participating in I/O interrupt processing and all CPUs other than the boot processor are capable of migrating while the vPar is running.

Prevention and Suggestions

In vPars releases prior to A.04 the only way to increase the I/O interrupt processing is to add additional bound CPUs. This does require a reboot of the vPar that is changing.

In vPars A.04 and A.05, all vPars start out able to handle I/O interrupt processing. However when CPUs are deleted online, interrupts are migrated out. But when CPUs are added online, interrupts are not automatically redistributed. With A.04 and A.05 you can issue the HP-UX command "inctl" to move interrupts from over-utilized CPUs to under-utilized CPUs.

Please see the whitepaper called "CPU Configuration Guidelines for vPars" for more details.

#8 – Memory Overhead with vPars

Symptoms: Performance, Functionality

The last vPar loaded on IPF systems has around 512MB of memory missing. Why is the last vPar short memory?

Prevention and Suggestions

On PA-RISC the vPar monitor and database simply reside in memory and take up around 20-25MB of memory one time regardless of the number of vPars. You will see the last vPar loaded short this 20-25MB.

On Itanium the vPar memory overhead is quite a bit more compared to PA-RISC. When the mode is changed to vPar, the vPar's monitor and supporting firmware needs a certain amount of memory to load itself and allocate its data structures. The total amount of memory taken is 384MB when the monitor is launched.

The memory required for each firmware instance associated with each vPar is typically around 32MB (and may be little more). And in addition, firmware takes away 64MB from each cell. Please note that this 64MB from each cell is also taken when in nPar mode and is not vPar specific.

So by the time you are done, the last vPar loaded may see as much as 512MB of missing. It will always be the last vPar loaded due to how the vPar monitor gives out memory as it loads the vPars in the order they are in the vPar database.

#9 – vPars and Securepath

Symptoms: Performance, Supportability

After installing Securepath the disk array LUNs under Securepath control will be claimed by hsx driver and it will no longer be possible to retrieve the EFI path to each device, i.e. neither 'ioscan -e' nor 'vparefiutil' will show EFI paths for these LUNs.

Prevention and Suggestions

When installing vPars A.04.01 on an IPF system which will also use Securepath the following installation procedure must be used:

- Install HP-UX

- Install PHKL_33581
- Install vPars
- vparcreate, install and boot individual partitions
- Make sure PHKL_33851 is installed on each partition
- Install Securepath in each partition

This order is necessary to overcome a Securepath limitation.

#10 – vPars and OLAR

Symptoms: Functionality

OLAR IS supported with vPars, just as it is on stand alone HP-UX systems.

Prevention and Suggestions

There are some things to keep in mind re OLAR and vPars, however:

- a) if an I/O path is not assigned to a virtual partition, neither SAM nor the OLAR commands can access it to operate on. Slots (even empty slots) that are going to be used for OLAR need to be assigned to a virtual partition.
- b) if the vPar that owns the slot is down, it must be booted so that SAM or the rad command can be run. In the case of a stand alone (non-vPars) system, if HP-UX is down the power to the hard partition may be turned off.

© 2006 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.

Itanium is a trademark or registered trademark of Intel Corporation in the U.S. and other countries and is used under license.

VERITAS, VERITAS Volume Manager, and VxVM are trademarks or registered trademarks of VERITAS Operating Corporation in the U.S. and other countries and are used under license.

Version 1.0, 02/2006