HP-UX 11i v3 Dynamic nPartitions

Features and configuration recommendations

Table of Contents

Introduction

This white paper introduces the Dynamic nPartitions feature that will debut with HP-UX 11i v3 Update 1. The nPartitions Background section summarizes the properties of nPartitions on HP-UX 11i before the update. The Dynamic nPartitions operations section explains the operations that make nPartitions dynamic. The Dynamic nPartitions commands section illustrates the invocation of the new Dynamic nPartitions commands. Finally, this paper concludes with Scenarios for Dynamic nPartitions and Configuration recommendations for Dynamic nPartitions to illustrate the value of Dynamic nPartitions and to show how to configure server complexes to realize the best value. Appendices cover an illustrated example of cell online migration, command sequences for Dynamic nPartitions scenarios, a transcript of a cell migration operation, troubleshooting hints, a glossary, and a table showing the servers that support Dynamic nPartitions.

nPartitions Background

Dynamic nPartitions functionality is an important extension of static nPartitions. Cell-based HP servers can be configured into multiple nPartitions of various sizes. The great value of cell-based servers is that they enable you to have incredible flexibility in configuration options. A server complex containing eight cells can be configured as a single nPartition containing all eight cells, eight independent single-cell nPartitions, or in any combination to suit the application workload. When the workload changes, the server complex can be repartitioned to contain a different number of independent nPartitions—each containing the optimal number of cells for its application. Before the introduction of Dynamic nPartitions operations, it was necessary to reboot a partition to change the cells contained in it. Dynamic nPartitions allows the composition of partitions to change online, while the nPartition continues to run its application workload, with no reboot or planned downtime needed.

For more information on nPartitions, refer to the nPartition Administrator's Guide. The appendix Cell-based HP servers and Dynamic nPartitions lists the HP servers that can be configured for nPartitions.

Dynamic nPartitions operations

There are two basic Dynamic nPartitions operations:

- Cell online activation
- Cell online deactivation

Cell online activation

The cell online activation operation is used to activate an inactive cell while the nPartition containing it continues to run. For this operation to be possible, an nPartition must be running HP-UX 11i Update 1 and its supporting firmware, and that nPartition must contain a cell in the inactive state. Upon completion of the cell online activation operation, the processor and memory resources on the activated cell are available for use by the operating system and applications, and any I/O resources will be ready to be activated.

Cell online deactivation

The cell online deactivation operation is used to deactivate an active cell while the nPartition containing it continues to run. For this operation to be possible, an nPartition must be running HP-UX 11i v3 Update 1 and its supporting firmware, that nPartition must contain a cell that is eligible for deactivation, and any I/O resources on that cell must be inactive. Upon completion of the cell online deactivation operation, the processor and memory resources on the deactivated cell will no longer be used by the operating system or applications.

Base cells and floating cells

A cell is eligible for online deactivation if and only if it is designated as a floating cell. Any cell not specifically designated as a floating cell is a base cell. HP-UX before 11i v3 included support for base cells only. A base cell cannot be online deactivated.

The system administrator designates a cell as a base cell or floating cell using the nPartition commands. Designation as a floating cell has two important implications. First, a floating cell can never contain any interleaved memory. 100% of the memory on a floating cell is cell local memory. Second, the operating system cannot place certain vital data elements in memory on floating cells. This precaution ensures that the operating system can continue to run at optimal performance if the floating cell is online deactivated.

The core cell for an nPartition is a base cell; therefore, every nPartition must contain at least one base cell. The system administrator might designate additional cells as base cells so that the operating system can have more memory in which to place its vital data elements.

A cell can be online activated as a base cell. In that case, it always contains 100% cell local memory, but it cannot subsequently be online deactivated.

A cell's designation as base or floating can be changed while the cell is active, however the change cannot take effect until the partition is rebooted. Similarly, the fraction of cell local memory on a base cell can be changed while the cell is active, but the change cannot take effect until the partition is rebooted.

Dynamic nPartitions commands

The static nPartition commands have been augmented with a new command, `parolrad`, to invoke cell online operations. This section gives the basics for the command. The full details are available in the manual pages.

The `parolrad` command can operate upon remote partitions, but all examples discussed here are for operations upon the local partition. That is, the partition running the operating system on which the `parolrad` is invoked is the target of the cell online operation.

The command for cell online activation is

```
parolrad –a cell_identifier
```

where *`cell_identifier`* specifies the cell to be online activated. The *`cell_identifier`* can be a global slot number or in the form cabinet number/slot number. It is required that the specified cell already be assigned to the local partition and be in the inactive state.

The command for cell online deactivation is

```
parolrad -d cell_identifier
```

where *`cell_identifier`* specifies the cell to be online deactivated. It is required that the specified cell be an active floating cell in the local partition.

Examples of the use of these commands are given in the appendix Illustrated example of online cell migration.

As an alternative to the command line interface, Dynamic nPartitions operations can also be initiated through the Partition Manager graphical user interface tool. In fact, the Partition Manager provides full support for all of the operations used to manage cell-based HP servers. Documentation for the tool is available on the Web.

I/O chassis

A cell can have an I/O chassis attached to it. Depending on the model of the server, a cell can also have core I/O associated with it.

If a cell with I/O (an I/O chassis or core I/O or both) is online activated, then the I/O is also available to the partition. However, the I/O is not automatically activated: a separate command, `olrad`, is required to activate the I/O.

The command for I/O chassis activation is

```
olrad –A -s cell_hardware_path
```

where the I/O chassis connected to the cell specified by *`cell_hardware_path`* is online activated. The *`cell_hardware_path`* must be a global slot number.

A cell with I/O can be online deactivated, but, for the operation to succeed, the I/O must first have been deactivated using the `olrad` command. The deactivation of I/O is subject to Critical Resource Analysis (CRA). If the CRA detects usages on the chassis to be deactivated, they will be noted in the logs.

The command for I/O chassis deactivation is

```
olrad –D -s cell_hardware_path
```

where the I/O chassis connected to the cell specified by *`cell_hardware_path`* is online deactivated. The *`cell_hardware_path`* must be a global slot number.

The olrad command can also be used to perform OLA, OLD, and OLR for PCI cards. However, such operations are intended to be performed when the I/O chassis is active. Changing PCI cards when the I/O chassis is inactive requires a power cycle sequence of the associated cell before the I/O chassis can be activated again. For more information, see the manual pages for `olrad`.

Scenarios for Dynamic nPartitions

Scenarios showing the value of Dynamic nPartitions are described in this section. Many more cases are possible.

Dynamic nPartitions functionality provides greater flexibility when more cells are in the system complex. A two-cell server does not allow a wide range of uses of Dynamic nPartitions. Accordingly, most examples are applicable to servers with four or more cells.

Adding capacity to an nPartition

Dynamic nPartitions functionality provides a means to add resources to an nPartition where the workload has increased beyond the original capacity of the partition. As long as the complex contains an unassigned cell, or a vacant slot into which a new cell can be inserted, the capacity of any nPartition can be increased with a cell online addition operation. Multiple cells can be added in succession, until all available cells in the entire complex are in use.

Shifting floating capacity to meet demand peaks

Dynamic nPartitions functionality is valuable when a server complex is divided into multiple partitions hosting application workloads that exhibit uncorrelated but predictable fluctuations in demand. The Dynamic nPartitions operations are fairly heavyweight, so the fluctuations must be of significant duration, a few hours or more. Different resource management tools are appropriate for workload spikes lasting only a few minutes or less.

For example, one partition hosts an application that is extremely busy reconciling financial accounts at the end of the month, while another partition does intensive data mining at the beginning of the next month. Alternatively, seasonal variations are likely when supporting retail operations, or applications might see extremely high loads only during the income tax preparation season. Finally, if different partitions supported call centers in different time zones, there could be regular, daily workload variations.

Figure 1 illustrates a configuration appropriate for this scenario. Suppose a server complex consisting of eight cells were partitioned with cells 0 and 1 as base cells in partition 0; cells 2 and 3 as base cells in partition 1; and cells 4 and 5 as base cells in partition 2. Cells 6 and 7 are floating cells that can be assigned to the partition that needs them most at the time.

Figure 1.

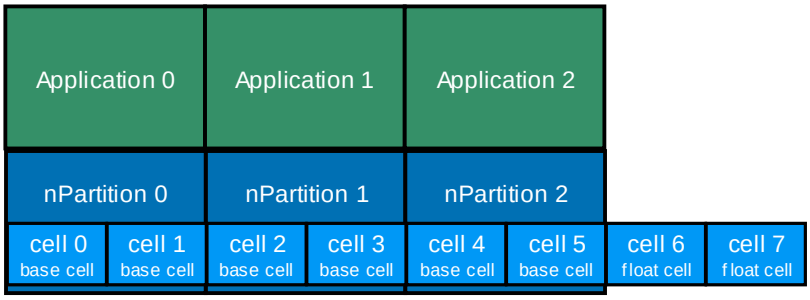

If the applications running on partition 0 experienced peak load when partitions 1 and 2 were relatively underutilized, cell online addition operations are used to add cells 6 and 7 to partition 0. If the workload pattern changed so that partition 2 needed more capacity, cell 6 is online deleted from partition 0 and then online added to partition 2. The two floating cells continue to be migrated among the partitions to follow the workload peaks.

A variation suggested by a customer is to place tape drives or other shared specialized peripherals on one of the floating cells. This floating cell could be migrated into the various partitions as need dictates. This would allow tape backups to be accomplished at local transfer speeds rather than network transfer speeds.

Temporary removal of a partition

Floating cells can be created by temporarily removing one or more partitions. For example, a server complex could be divided into a production partition and a development and test partition. During times of urgent production need, such as a retailer satisfying a holiday peak, you can shut down the development and test partition and use its cells as floating cells to be online added to the production partition. When the production peak has abated, the floating cells can be online deleted to reconstitute the development and test partition.

Replacement of hardware

In the preceding scenarios, cells migrated from one partition to another while remaining within the cellular complex. When a cell is online deleted from a partition, it can then be powered down and physically removed from its cabinet. This would allow components on the cell board or the entire cell board to be replaced.

The replacement capability enables hardware maintenance to be performed, for example, if predictive diagnostics indicated that a certain memory DIMM or processor module was likely to fail and should be replaced. It is possible to increase the number of DIMMs in a cell board or replace DIMMs with denser components.

Only floating cells are available for this type of replacement. You must reboot the partition to replace a base cell. Also, this is not the same as the "online replacement" available for PCI cards, which involves a temporary suspension of access and replacement with identical hardware configured into the identical state.

Configuration recommendations for Dynamic nPartitions

This section contains guidelines for optimizing the benefits of Dynamic nPartitions.

Minimum amount of memory on base cells

The operating system never places certain vital data elements in memory on floating cells. This is an important optimization to keep the operating system running at the highest possible performance level. Therefore, the system administrator must configure each partition with a sufficient number of base cells, or, more specifically, a sufficient amount of memory on base cells to meet the operating system requirements.

The following table describes the rule for the total amount of physical memory present in the partition as compared to the minimum amount of memory on base cells.

Physical memory (total in the partition)	Minimum amount of memory on base cells
1 GB to 8 GB	1/2 of the physical memory
8 GB to 16 GB	4 GB
16 GB or more	1/4 of the physical memory

If, as is typical, the same amount of memory is on each cell, with that amount being 8 GB or more, the rule can be stated simply: configure at least one-quarter of the cells in a partition as base cells.

The ratio of the amount of memory on base cells to the amount of memory on floating cells can have significant performance implications. A partition with insufficient memory to place vital kernel data might not perform well. An important special case is the size of the HP Online JFS buffer cache. If the "autotune" option for the size of the HP Online JFS buffer cache is selected, the cache will be automatically resized upon Dynamic nPartitions operations. The consequence is that the buffer cache may be smaller on a system with many floating cells than on a comparable system comprised only of base cells. The smaller buffer cache may reduce the performance of certain file system operations, such as "`find`".

The fact that floating cells must contain 100% cell local memory also influences the choice to configure a cell as base or floating. For workloads with strong and predictable memory locality, the use of cell local memory and process-to-processor affinity can yield significant performance improvements. Workloads that exhibit no memory locality may perform better with more interleaved memory.

I/O chassis

Many workloads are CPU-bound or memory-bound. When using cell online activation to increase the capacity of a partition running such a workload, an I/O chassis does not need to be attached to the cell.

In some cases, I/O resources can be considered private to an application and to the partition that hosts it. An example would be a SCSI controller that connects to disks containing application data. In such cases, it is not appropriate for the I/O resources to migrate to another partition along with a floating cell.

The configuration guideline then is to configure the cellular complex so that private I/O resources are attached to base cells. Floating cells used to satisfy demand for CPU and memory resources do not need an I/O chassis attached. Any I/O attached to a floating cell must be shareable, such as additional redundant connections to the network or to the storage fabric.

Duration of Dynamic nPartitions operations

The cell online activation operation must configure all resources on the new cell for use by the operating system. It also rebalances some operating system parameters to suit the larger size of the nPartition. For this reason, the operation can take some time to complete. The typical time to complete cell online activation on a moderately loaded system is 60 seconds.

The cell online deactivation operation causes the operating system to discontinue use of all resources on the cell to be deactivated. In particular, all data elements in memory must be preserved, either in other memory or on disk. Some operating system parameters are rebalanced to suit the smaller size of the nPartition. Therefore, the operation can take some time to complete. The typical time to complete cell online deactivation on a moderately loaded system is 3 minutes.

If a Dynamic nPartitions operation is in progress, its status can be ascertained by the command

```
parolrad –m
```

If a Dynamic nPartitions operation is in progress and has not yet reached its commit point, it can be canceled by the command

```
parolrad –c sequence_number
```

where `sequence_number` is the number returned by the invocation of the parolrad command that initiated the Dynamic nPartitions operation. Most Dynamic nPartitions operations reach their commit point within a few seconds and then can no longer be canceled.

Sensitivity to operating system workload

The duration of a Dynamic nPartitions operation is sensitive to the workload on the target partition. A cell online activation operation is often used to increase the resources in a partition when it is inadequate to meet the demands of the workload. Therefore, it is usual and expected to perform cell online activation to a partition that is overutilized, and the duration of such an operation is not expected to exceed 5 minutes.

By contrast, a cell online deactivation operation decreases the resources in a partition. So, it should generally not be performed on a partition that is already overutilized. An attempt to delete a cell from a partition with a high memory utilization level could take as long as 1 hour.

If the system load average is extremely high, that is, above 40, the parolrad command may time out without initiating the operation. In that case, it would be appropriate to repeat the command, because it will typically succeed the second time.

The system administrator should reduce the workload in the partition to a level that can be sustained by the remaining resources before performing a cell online deactivation operation. The workload can be reduced by redirecting network traffic to other servers, by reducing the number of interactive user

sessions, or similar measures. Various capacity planning tools are available to match workloads to hardware resource levels.

These considerations apply particularly to the Oracle Database products. If the memory allocated to the Oracle Database is based on the full complement of cells in the partition, then the system will not perform well if cells were to be deactivated. Instead, it would be appropriate to use the Oracle tuning mechanisms, for example, the sga_target parameter, to make the memory footprint of the database consistent with the size of the memory that will remain after cells are deactivated.

Current limits and conditions

Rollout of Dynamic nPartitions functionality in HP-UX 11i v3

Dynamic nPartitions operations will be available with HP-UX 11i v3 Update 1. HP-UX 11i v3 will run successfully on servers that predate the release of the update. However, the new Dynamic nPartitions operations will not be functional unless the firmware on the server is upgraded to the firmware that supports the new features. The minimum firmware version for each platform that supports Dynamic nPartitions is shown in the Firmware Matrix.

Dynamic nPartitions, virtual partitions, and Integrity Virtual Machines

HP-UX 11i v3 Virtual Partitions enable dynamic operation through processor and memory migration among vPars. However, Dynamic nPartitions operations cannot be performed to an nPartition that is running vPars.

Integrity Virtual Machines also enable dynamic operation through processor and memory migration among guest instances. However, the Integrity Virtual Machines Platform Manager as of September 2007 is based on HP-UX 11i v2. Therefore, Dynamic nPartitions operations cannot be performed to an nPartition that is running Integrity Virtual Machines.

Requirement to reset Workload Manager

Workload Manager (WLM) exercises control over the hardware resources in a partition or a set of partitions. WLM uses its own algorithms to adjust hardware resources to meet the stipulated service levels. Those algorithms are perturbed if the system administrator uses Dynamic nPartitions to change the composition of partitions being managed by WLM. Therefore, in cases where WLM is being used, it is necessary to shut down the tool before a Dynamic nPartitions operation and restart the tool after the operation completes.

The situation is similar with the global Workload Manager (gWLM). gWLM expects to have complete control of the CPUs, or cores, available on a system. However, when you must make manual adjustments to system resources, you can do so as follows:

1. Undeploy the SRD containing the systems that you want to adjust
2. Perform the Dynamic nPartitions operations
3. Re-create and re-deploy the SRD

Integration with instant capacity

Dynamic nPartitions is fully integrated with instant capacity functionality. Any processor licensing restrictions imposed by instant capacity are respected by Dynamic nPartitions. For that reason, a cell online activation operation is possible only if there is at least one processor license available in the cellular complex.

Interaction with processor sets

Dynamic nPartitions is fully integrated with processor sets. If the attributes of a processor set are established so that it is not allowed to become empty, then the last SPU in that processor set cannot be online deactivated. If a cell online deactivation operation would cause the processor set to become empty, then that operation will not complete successfully.

Cell reset recovery time

When a cell is deactivated online, it is reset to prepare for its next use. The reset operation does a full initialization, so it takes approximately 1 minute. During this time the cell can not be reactivated nor assigned to another partition. During the reset recovery time, parstatus shows the state of the cell as "Powering on".

Summary

Dynamic nPartitions operations complete the value proposition of partitionable server complexes. Vast flexibility can be realized through changing the composition of nPartitions while continuing to run their application workloads. Many real-world scenarios can benefit from timely use of Dynamic nPartitions. Advance planning in the configuration of the system complex can maximize the benefits of Dynamic nPartitions functionality.

Illustrated example of online cell migration

This example of using Dynamic nPartitions operations to migrate a cell from one partition to another features a server complex containing eight cells. The complex is partitioned with cells 0, 1, and 2 in nPartition 0. Cells 3, 4, 5, 6, and 7 are in nPartition 1. Each partition runs a separate instance of the HP-UX operating system. Figure 2 shows the initial configuration.

Figure 2.

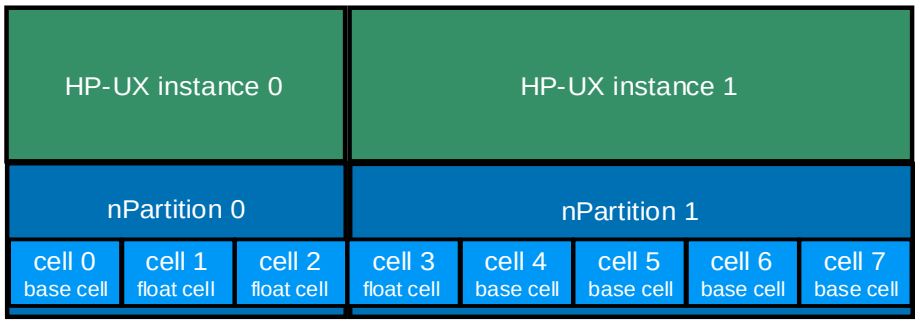

If partition 0 is underutilized, or if partition 1 is running a higher priority workload, the system administrator might decide to migrate cell 2 from partition 0 to partition 1. Assuming that cell 2 does not have associated I/O, the migration sequence begins with the command

```
parolrad -d 2
```

This command, entered in a terminal session on HP-UX instance 0, causes HP-UX to stop using all resources on cell 2, and it causes the cell to be reset. The new state of the complex is shown in Figure 3:

Figure 3.

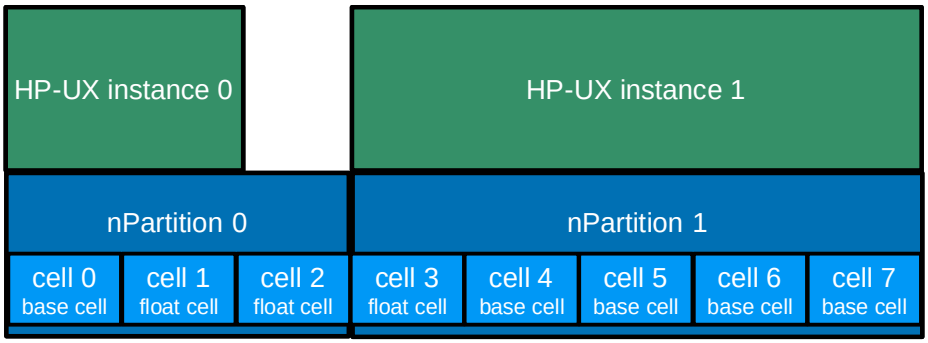

Cell 2 is still assigned to partition 0, but it is inactive. Now, the cell can be unassigned from partition 0 with the command

```
parmodify –p 0 -d 2
```

Figure 4.

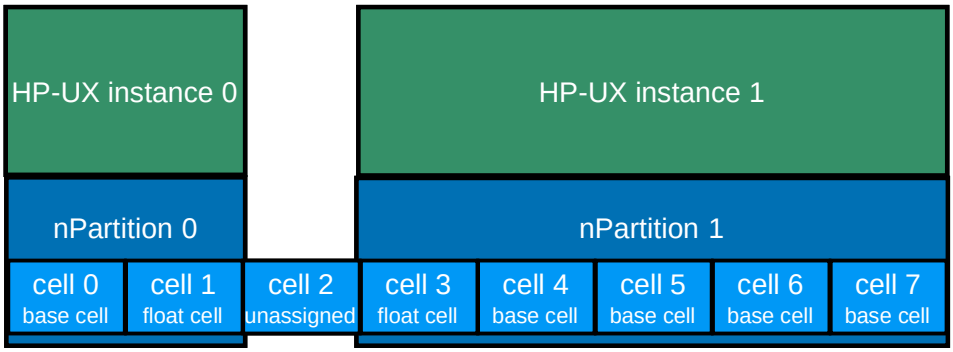

Cell 2 is not assigned to any partition.
For this example, the system administrator does not use the remote management capabilities of the nPartition commands. Instead, the following commands are entered into a terminal session on HP-UX instance 1.

Cell 2 is unassigned, and once the cell reset operation triggered by the previous online deactivation has completed, it can be assigned to partition 1 with the command

```
parmodify –p 1 -a 2:floating:y::100%
```

The attributes designate cell 2 as a floating cell with 100% of its memory used as cell local memory. The result of the assignment is shown in the following diagram:

Figure 5.

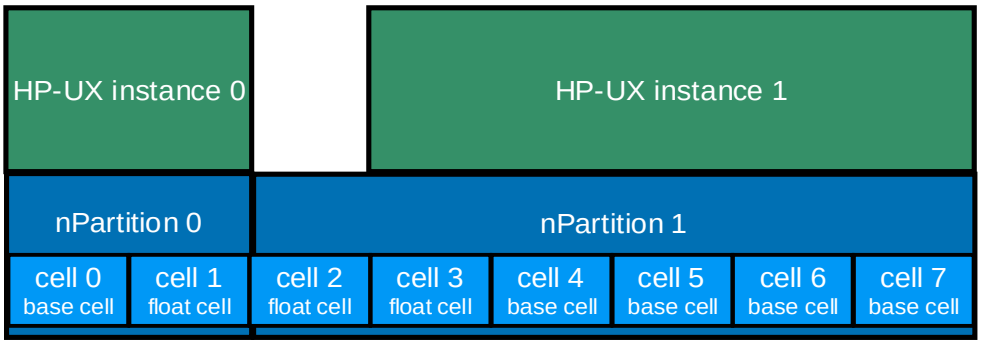

Note that cell 2 is assigned to partition 1, but it is inactive, so its resources are not being used by HP-UX instance 1. The cell can be online activated with the command

```
parolrad -a 2
```

Figure 6.

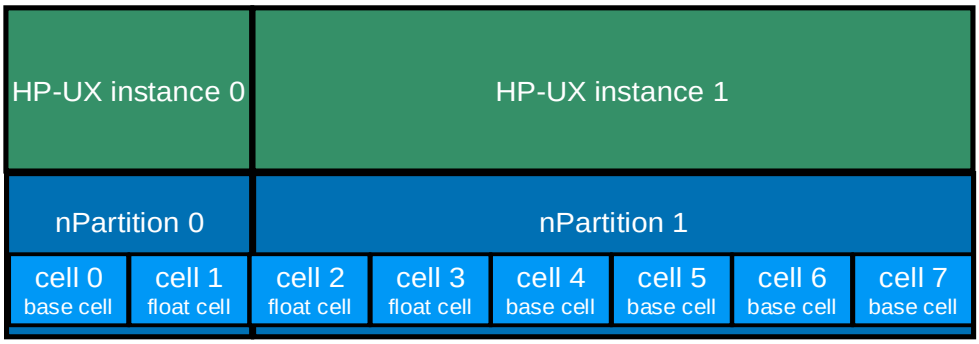

Now, the sequence is complete: cell 2 is fully integrated into HP-UX instance 1. Both operating systems continued to run their application workload throughout the entire sequence.

Command sequences for Dynamic nPartitions scenarios

This section shows representative command sequences to perform common Dynamic nPartitions operations.

Cell replacement

The following command sequence could be used if it is desired to replace or perform hardware maintenance upon the cell whose global slot number is **S**. It is assumed that the target cell is a floating cell and it has an active I/O chassis connected to it.
First, it is necessary to quiesce all activity associated with the I/O chassis connected to the target cell. If there are any active LAN cards, it is necessary to use the ifconfig command with the unplumb option to tear down the Streams plumbing. If there are filesystems mounted on devices connected through the I/O chassis, then it is necessary to unmount them. The olrad command will give a CRA report to indicate any I/O usages if the I/O chassis to be deactivated is not quiescent.
Once the I/O is quiesced, the cell replacement sequence can be executed:

```
olrad -D -s S
parolrad -d S
frupower -f -c S
```

The cell with global slot number **S** can now be physically removed from the cabinet. The cell can be repaired and reinserted, or it can be replaced with a new compatible cell.

```
frupower -o -c S
```

It is necessary to wait until the cell has completed its power on sequence. The parstatus command will show the cell in the "Inactive" state when the cell is ready.

```
parolrad -a S
olrad -A -s S
```

Now the cell in the I/O connected to it are activated. The LAN connections can be reestablished in the filesystems remounted.
Note that it was not necessary to use the parmodify command to unassign the cell from its partition.

NOTE: During the Cell replacement sequence as mentioned above, when the I/O chassis is activated (using the "olrad -A -s" command), the instance numbers that get assigned for the I/O devices in the Chassis may be different from the instance numbers that were assigned for the I/O devices before the Chassis was deactivated (using the "olrad -D -s" command). This may result in I/O devices having different device special files(dsf) after a Cell replacement operation (persistent device special files are not impacted). If the user needs to have the same instance numbers (and thereby same dsf) for the I/O devices after the Cell Replacement sequence is performed, the "-f <file> -r " option of ioinit(1M) can be used to reassign the previous instance numbers. Note that the “–r” option must be specified with ioinit(1M) as some of the I/O devices of type ext_bus, lan do not support online instance number reassignment.

Cell migration

Here is the command sequence to migrate the cell in slot **S** from partition **p** to partition **q**. It is assumed that the target cell is a floating cell and it does not have an I/O chassis. (If it did have active I/O, the instructions above for quiesceing and reestablishing I/O would apply.) Further, it is assumed that security certificates are in place to enable remote invocation of the parolrad command.

The system name of the HP-UX instance running in partition **q** is "remote_system" and the userid "administrator" has the required privilege.

```
parolrad -d S
parmodify -p p -d S
parmodify -p q -a S:floating:y:float:100%
```

Is necessary to wait until the cell has completed its power on sequence.

```
parolrad -a S -u administrator: -g remote_system
```

The parolrad command will prompt for the password associated with the administrator userid. Then, the cell migration operation will complete.

Transcript of cell migration operation

Following is an actual transcript of a cell migration operation, performed on a system with a prototype version of the Dynamic nPartitions feature. The objective of this operation is to migrate cell 2 from partition 1 to partition 0.
The first commands are typed to a terminal session on partition 1.

```
[par1]# parstatus -C

[Cell]
                        CPU     Memory                                  Use
                        OK/     (GB)                            Core    On
Hardware   Actual       Deconf/ OK/                             Cell    Next Par
Location   Usage        Max     Deconf    Connected To          Capable Boot Num
========== ============ ======= ========= ==================== ======= ==== ===
cab0,cell0 Active Core  8/0/8   8.0/0.0   cab0,bay0,chassis0  yes     yes  0
cab0,cell1 Active Core  8/0/8   8.0/0.0   cab0,bay0,chassis1  yes     yes  1
cab0,cell2 Active Float 8/0/8   8.0/0.0   cab8,bay0,chassis0  yes     yes  1
cab0,cell3 Active Base  8/0/8   8.0/0.0   cab8,bay0,chassis1  yes     yes  0

Notes: * = Cell has no interleaved memory.
[par1]# olrad -D -s 2

Chassis Delete operation in progress...
[ NOTE: The chassis delete operation may take a few minutes to complete
on large configurations. It is recommended not to disrupt this operation. ]

CRA REPORT SUMMARY: CRA returned SUCCESS.

Detailed CRA report is available in /var/adm/cra.log file.

Activity      : End of Chassis-Delete operation. Ok to go ahead with the
                OnLine De-activation of the Cell.
Hardware path : 2

[par1]# parolrad -d 2
Note: Cell Online deactivation operation has been initiated.
The sequence number is: 2

Note: Cell OL* operations, may take several minutes to complete.
Killing or otherwise aborting this parolrad command will not cancel the
Cell OL* operation. See the parolrad (1M) manpage for details.

Online operation in progress ....

Online operation has completed successfully.
[par1]# parmodify -p 1 -d 2
[par1]# parstatus -C
```

```
[Cell]
                        CPU     Memory                                  Use
                        OK/     (GB)                            Core    On
Hardware   Actual       Deconf/ OK/                             Cell    Next Par
Location   Usage        Max     Deconf    Connected To          Capable Boot Num
========== ============ ======= ========= ==================== ======= ==== ===
cab0,cell0 Active Core  8/0/8   8.0/0.0   cab0,bay0,chassis0    yes     yes  0
cab0,cell1 Active Core  8/0/8   8.0/0.0   cab0,bay0,chassis1    yes     yes  1
cab0,cell2 Inactive     8/0/8   8.0/0.0   cab8,bay0,chassis0    yes     -    -
cab0,cell3 Active Base  8/0/8   8.0/0.0   cab8,bay0,chassis1    yes     yes  0

Notes: * = Cell has no interleaved memory.
```

Note that the output from the parstatus command indicates that cell 2 has completed its power on reset sequence and is in the "Inactive" state.
The demonstration system was not configured for remote invocation of parolrad, so the remaining commands are issued in partition 0.

```
[par0]# parstatus -C

[Cell]
                        CPU     Memory                                  Use
                        OK/     (GB)                            Core    On
Hardware   Actual       Deconf/ OK/                             Cell    Next Par
Location   Usage        Max     Deconf    Connected To          Capable Boot Num
========== ============ ======= ========= ==================== ======= ==== ===
cab0,cell0 Active Core  8/0/8   8.0/0.0   cab0,bay0,chassis0    yes     yes  0
cab0,cell1 Active Core  8/0/8   8.0/0.0   cab0,bay0,chassis1    yes     yes  1
cab0,cell2 Inactive     8/0/8   8.0/0.0   cab8,bay0,chassis0    yes     -    -
cab0,cell3 Active Base  8/0/8   8.0/0.0   cab8,bay0,chassis1    yes     yes  0

Notes: * = Cell has no interleaved memory.
 [par0]# /locinfo2
 --- System wide locality info: ---
 index  ldom physid  type     total      free     float      used
     0     0      0   CLM     3.93G     3.36G         0   587.36M
     1     1      3   CLM     3.94G     3.54G         0   406.21M
     2    -1     -1   ILV     8.00G     6.26G         0     1.74G
                             -----     -----     -----     -----
                           15.87G    13.16G         0     2.71G

[par0]# parmodify -p 0 -a 2:floating:y:float:100%
Note: It may take a longer time for the partition to boot if
any cell is performing power-on selftest.

To activate newly added cells, reboot the partition for reconfiguration.
or activate a Cell Online using "parolrad" command.
On HP-UX use "shutdown -R" or "parolrad -a".
On Windows, restart the OS on the partition using the Start menu or "shutdown -r
" at a Command Prompt.
On Linux, restart the OS on the partition using "shutdown -r" at a Command Promp
t.
If the partition is at the system firmware prompt, use the RR command on the MP.

[par0]# parolrad -a 2
Note: Cell Online activation operation has been initiated.
The sequence number is: 3

Note: Cell OL* operations, may take several minutes to complete.
Killing or otherwise aborting this parolrad command will not cancel the
Cell OL* operation. See the parolrad (1M) manpage for details.

Online operation in progress ....
```

```
Online operation has completed successfully.
[par0]# olrad -A -s 2

Chassis Add operation in progress...
[ NOTE: The chassis add operation may take a few minutes to complete
on large configurations. Killing or otherwise aborting this olrad
command may not cancel the Chassis Add operation. ]

H/W Path   Class                  Description
============================================
2          cell
2/0            ioa                System Bus Adapter (12eb)
2/0/0             ba              Local PCI-X Bus Adapter (122e)
2/0/0/1/0               lan       HP PCI-X 1000Base-T Built-in
2/0/0/2/0               ext_bus   SCSI Ultra320
2/0/0/2/1               ext_bus   SCSI Ultra320
2/0/0/3/0               ext_bus   SCSI Ultra320
2/0/0/3/1               ext_bus   SCSI Ultra320
2/0/1             ba              Local PCI-X Bus Adapter (12ee)
2/0/1/1              slot         PCI Slot
2/0/2             ba              Local PCI-X Bus Adapter (12ee)
2/0/2/1              slot         PCI Slot
2/0/4             ba              Local PCI-X Bus Adapter (12ee)
2/0/4/1              slot         PCI Slot
2/0/6             ba              Local PCI-X Bus Adapter (12ee)
2/0/6/1              slot         PCI Slot
2/0/8             ba              Local PCI-X Bus Adapter (12ee)
2/0/8/1              slot         PCI Slot
2/0/10            ba              Local PCI-X Bus Adapter (12ee)
2/0/10/1             slot         PCI Slot
2/0/12            ba              Local PCI-X Bus Adapter (12ee)
2/0/12/1             slot         PCI Slot
2/0/14            ba              Local PCI-X Bus Adapter (12ee)
2/0/14/1             slot         PCI Slot
2/120          processor          Processor
2/121          processor          Processor
2/122          processor          Processor
2/123          processor          Processor
2/124          processor          Processor
2/125          processor          Processor
2/126          processor          Processor
2/127          processor          Processor
2/250          ba                 Core I/O Adapter
2/250/0           acpi_node       Acpi Hardware


SCSI port at hardware path, 2/0/0/2/0, configured.
Installing special file for MPT instance 12
SCSI port at hardware path, 2/0/0/2/1, configured.
Installing special file for MPT instance 13
SCSI port at hardware path, 2/0/0/3/0, configured.
Installing special file for MPT instance 14
SCSI port at hardware path, 2/0/0/3/1, configured.
Installing special file for MPT instance 15

[par0]# parstatus -C

[Cell]
                       CPU     Memory                                  Use
                       OK/     (GB)                            Core    On
Hardware   Actual      Deconf/ OK/                             Cell    Next Par
Location   Usage       Max     Deconf    Connected To          Capable Boot Num
========== =========== ======= ========= ==================== ======= ==== ===
cab0,cell0 Active Core  8/0/8   8.0/0.0   cab0,bay0,chassis0  yes     yes  0
cab0,cell1 Active Core  8/0/8   8.0/0.0   cab0,bay0,chassis1  yes     yes  1
```

```
cab0,cell2 Active Float 8/0/8   8.0/0.0   cab8,bay0,chassis0  yes     yes  0
cab0,cell3 Active Base  8/0/8   8.0/0.0   cab8,bay0,chassis1  yes     yes  0

Notes: * = Cell has no interleaved memory.
[par0]# /locinfo2
 --- System wide locality info: ---
 index  ldom physid  type     total      free     float      used
     0     0      0   CLM     3.93G     3.36G         0   587.37M
     1     1      3   CLM     3.94G     3.53G         0   412.05M
     2     2      2   CLM     7.94G     7.55G     7.55G   398.78M
     3    -1     -1   ILV     8.00G     6.24G         0     1.76G
                             -----     -----     -----     -----
                           23.81G    20.68G     7.55G     3.13G
```

What could possibly go wrong?

Dynamic nPartitions commands return a message if the operation cannot proceed as intended. The following table gives some troubleshooting hints.

Message	Likely Cause	Remedy
The cell shall be assigned to the local (target) nPar.	The cell to be online activated or deactivated is not already a member of the target partition.	Assign the cell to the partition before activating it.
The cell shall be at BIB (completed POST).	In the case of online cell migration, the cell to be online activated is still completing its power on initialization after being online deactivated from its previous partition.	Wait for one minute, or until parstatus shows the state of the target cell to be "Inactive".
The Use On Next Boot flag shall be set to true for the target cell.	The Use On Next Boot flag was set to false when the target cell was added to its partition.	Use the parmodify command to set the Use On Next Boot flag to true.
There shall not be a cell OL* operation already in progress on the local (target) nPar.	There is already a cell online activation or deactivation operation in progress in the target partition.	Wait for the previous operation to complete. The command parolrad -m will monitor the progress of that operation.
The OS on the target nPar shall support cell online operations.	The Dynamic nPartitions feature is not installed.	Install the DynamicNPars product from HP-UX 11i v3 Update 1.
vPar monitor shall not be running on the target nPar.	The parolrad command was issued from a vPars instance.	Use the vPars processor and memory migration facilities instead of parolrad.
The system FW in the active cells in the target nPar shall support cell OLAD.	The firmware on the cells already in the partition does not support Dynamic nPartitions.	Upgrade the system firmware on all cells to at least the minimum version specified in the firmware matrix.
The system FW in the target cell shall support cell OLAD.	The firmware on the cell being activated does not support Dynamic nPartitions.	
The target cell shall be compatible with the existing active cells in the local (target) nPar.	The cell being activated is not compatible with the cells already in the partition.	Make adjustments so that all of the cells have the same processor type and the same system firmware version.
There shall not be a pending profile at the time of the OL* operation initiation.	An administrator has used the parmodify command to modify the partition in a way that requires a reboot for the change to take effect.	Reboot the partition so that the pending profile can take effect.
The cell shall be active.	You are trying to deactivate a cell that is not already active.	Check for a typo in the cell identifier.
The cell shall be a floating cell.	You are trying to deactivate a base cell.	Check for a typo in the cell identifier.
All cards in the I/O chassis attached to the target cell shall be deconfigured.	The cell been deactivated has an active I/O chassis.	Quiesce the I/O, then use the olrad command to deactivate the I/O chassis.

Glossary

Term	Definition
Assigned and unassigned cells	Each cell in a server complex either is assigned to one of the nPartitions in the complex, or it is unassigned and thus is not used by any of the nPartitions. If an I/O chassis is attached to an unassigned cell, then the chassis likewise is not assigned to an nPartition. Cells that are unassigned are considered to be available resources; they are free to be assigned to any of the existing nPartitions, or can be used to create new nPartitions.
Active and inactive cells	Cells that are assigned to an nPartition and have booted to form an nPartition are active cells whose resources (processors, memory, and any attached I/O) can be actively used by software running in the nPartition. Cells that are inactive either are not assigned to an nPartition, or they have not participated in partition rendezvous to form an nPartition with any other cells assigned to the nPartition. (Partition rendezvous is the point during the nPartition boot process when all available cells in an nPartition join together to establish which cells are active for the current boot of the nPartition.)
Cell compatibility rules	While cells of different types can be freely mixed within a system complex, only compatible cells can be included in the same nPartition. The exact compatibility rules vary with the server model, but generally the compatibility rules restrict partitions to contain cells using the same processor type. The cell compatibility rules apply whether the nPartition is configured statically or through Dynamic nPartitions operations.
Online cell addition	An unassigned cell can be added to a partition (with the parmodify command) and immediately online activated (with the parolrad command). The combination of these two steps is called online cell addition.
Online cell deletion	A floating cell can be deactivated from a partition (with the parolrad command) and then immediately unassigned from the partition (with the parmodify command). The combination of these two steps is called online cell deletion.
Online cell migration	Online cell migration refers to the online deletion of a cell from one partition and its online addition to another partition.
Physical cell insertion	Physical cell insertion refers to the insertion of a cell board into a previously vacant cell slot in a server complex. The inserted cell becomes unassigned. None of the partitions in the complex are affected by this operation.
Physical cell removal	Physical cell removal refers to the removal of an inactive (and unpowered) cell board from a server complex. None of the partitions in the complex are affected by this operation.
Online cell replacement	Replacement of a cell board by the sequence of operations: online cell deletion, physical cell removal, physical cell insertion of a new cell into the same slot, followed by online cell addition.

Cell-based HP servers and Dynamic nPartitions

Cell-based HP servers that support HP-UX 11i v3 Dynamic nPartitions operations:

Cell capacity	Server model
Two-cell Servers	HP 9000 rp7420 Server
	HP 9000 rp7440 Server
	HP Integrity rx7620 Server
	HP Integrity rx7640 Server
Four-cell Servers	HP 9000 rp8420 Server
	HP 9000 rp8440 Server
	HP Integrity rx8620 Server
	HP Integrity rx8640 Server
HP 9000 Superdome Servers	HP 9000 Superdome Server (PA 8800 and PA 8900)
	HP 9000 Superdome Server with sx2000
HP Integrity Superdome Servers	HP Integrity Superdome Server
	HP Integrity Superdome Server with sx2000

Cell-based HP servers that do not support HP-UX 11i v3 Dynamic nPartitions operations:

Cell capacity	Server model
Two-cell Servers	HP 9000 rp7405/7410 Server
Four-cell Servers	HP 9000 rp8400 Server
HP 9000 Superdome Servers	HP 9000 Superdome Server (PA 8600 and PA 8700)

For more information

http://www.hp.com/go/hpux11iv3

http://www.docs.hp.com/en/5991-1247B/ for the nPartition Administrator's Guide.

http://docs.hp.com/en/PARMGR2/ for a description of the Partition Manager product.

http://docs.hp.com/en/SFWM1/index.html for the firmware versions to support Dynamic nPartitions.

4AA0-10907, February 2008