

Power Systems

Live Partition Mobility

IBM

Power Systems

Live Partition Mobility

IBM

Note

Before using this information and the product it supports, read the information in "Notices" on page 145.

This edition applies to IBM Virtual I/O Server Version 2.2.3.2 and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2012, 2014.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Live Partition Mobility	1
What's new in Live Partition Mobility	1
Live Partition Mobility on HMC-managed systems	4
Partition mobility overview for HMC	4
Benefits of partition mobility	4
Partition mobility process	4
Configuration validation for partition mobility	6
Logical partition attributes that change after the logical partition moves to the destination system	11
Processor compatibility modes	12
Processor compatibility mode definitions	12
Current and preferred processor compatibility modes	14
Enhanced processor compatibility modes	17
Migration combinations of processor compatibility modes	17
Examples: Using processor compatibility modes in partition mobility	27
Partition mobility environment	29
Source and destination servers in a partition mobility environment	29
Hardware Management Console in a partition mobility environment	30
Source and destination Virtual I/O Server logical partitions in a partition mobility environment	31
Live Partition Mobility pseudodevice	37
Configuring the VIOS firewall for partition mobility	40
Mobile partition managed by an HMC in a partition mobility environment	41
Software applications that recognize partition mobility	42
Network configuration in a partition mobility environment	42
Storage configuration in a partition mobility environment	43
Preparing for partition mobility	48
HMC-managed systems: Preparing the source and destination servers for partition mobility	48
Determining the available physical memory on the destination server	53
Determining the available I/O entitled memory on the destination server	53
Defining the partition profile policy for inactive partition mobility	55
Verifying that the destination server supports suspend-capable partitions	55
Determining the reserved storage device size in the destination server	55
Verifying the processor-level hardware capabilities of the destination server	55
Determining available processors on the destination server	56
Improving the performance of partition mobility	56
Server evacuation	57
Preparing the HMC for partition mobility	58
Verifying SSH authentication between the source and destination HMCs	58
Preparing the source and destination Virtual I/O Server logical partitions for partition mobility	59
Enabling the source and destination mover service partitions	61
Verifying that the destination shared memory pool contains an available paging space device	62
Preparing the mobile partition for partition mobility	63
Verifying RMC connections for the mobile partition	65
Verifying the processor compatibility mode of the mobile partition	66
Disabling the mobile partition for redundant error-path reporting	67
Disabling virtual serial adapters for the mobile partition	67
Removing the mobile partition from a partition workload group	68
Disabling BSR arrays for the mobile partition	68
Disabling huge pages for the mobile partition	69
Removing logical Host Ethernet Adapters from the mobile partition	70
Preparing the network configuration for partition mobility	71
Configuring secure IP tunnels between the mover service partitions on the source and destination servers	72
Preparing the virtual SCSI configuration for partition mobility	74
Setting the reserve policy attributes of a device	75
Verifying the virtual adapter connections between the mobile partition and the Virtual I/O Server logical partitions on the source server	76
Verifying that the mobile partition has access to its physical storage	77

Specifying a new name for a virtual target device to use on a destination VIOS partition	78
Preparing the virtual Fiber Channel configuration for partition mobility	79
Identifying the WWPNs that are assigned to a virtual Fibre Channel adapter	80
Verifying the virtual adapter connections between the mobile partition and the Virtual I/O Server logical partitions on the source server	81
Validating the configuration for partition mobility	82
Moving the mobile partition	83
Moving the mobile partition with HMC	83
Adding the mobile partition to a partition workload group	85
Moving the suspended mobile partition with the HMC command-line interface	86
Resuming the suspended mobile partition with HMC	86
Shutting down the suspended mobile partition with HMC	87
Troubleshooting partition mobility	87
Troubleshooting active partition mobility	87
Troubleshooting inactive partition mobility	92
Virtual I/O Server errors	93
Live Partition Mobility on IVM-managed systems	94
Partition mobility overview for IVM	94
Benefits of partition mobility	94
Partition mobility process for IVM.	94
Configuration validation for partition mobility	96
Logical partition attributes that change after the logical partition moves to the destination system	99
Processor compatibility modes	99
Processor compatibility mode definitions	100
Current and preferred processor compatibility modes	102
Enhanced processor compatibility modes	104
Migration combinations of processor compatibility modes	105
Examples: Using processor compatibility modes in partition mobility.	115
Partition mobility environment	116
Source and destination servers in a partition mobility environment	117
Integrated Virtualization Manager in a partition mobility environment	117
Software applications that recognize partition mobility	119
Network configuration in a partition mobility environment	119
Storage configuration in a partition mobility environment	120
Preparing for partition mobility	123
IVM-managed systems: Preparing the source and destination servers for partition mobility	123
Determining the available physical memory on the destination server	126
Determining the available I/O entitled memory on the destination server	126
Verifying that the destination server supports the virtual server network	127
Determining the virtual Ethernet switch name and mode in the destination server	128
Determining available processors on the destination server	128
Preparing the source and destination management partitions for partition mobility	129
Verifying that the destination shared memory pool contains an available paging space device	130
Preparing the mobile partition for partition mobility	131
Configuring the Virtual I/O Server for the VSN capability	132
Verifying the processor compatibility mode of the mobile partition	133
Removing the mobile partition from a partition workload group	134
Preparing the network configuration for partition mobility	134
Configuring secure IP tunnels between the mover service partitions on the source and destination servers	135
Preparing the virtual SCSI configuration for partition mobility	137
Setting the reserve policy attributes of a device	137
Verifying that the mobile partition has access to its physical storage	139
Specifying a new name for a virtual target device to use on a destination management partition	139
Preparing the virtual Fiber Channel configuration for partition mobility	140
Verifying the number of physical Fibre Channel ports that are available on the destination management partition	141
Validating the configuration for partition mobility	142
Moving the mobile partition	142

Notices	145
Programming interface information	147
Trademarks	147
Terms and conditions.	147

Live Partition Mobility

Live Partition Mobility, a component of the PowerVM[®] hardware feature for IBM PowerLinux servers, provides the ability to move Linux logical partitions from one system to another. The mobility process transfers the system environment that include the processor state, memory, attached virtual devices, and connected users.

Active partition mobility allows you to move Linux logical partitions that are running, including the operating system and applications, from one system to another. The logical partition and the applications running on that migrated logical partition do not need to be shut down.

Inactive partition mobility allows you to move a powered off Linux logical partition from one system to another.

You can use the Hardware Management Console (HMC), the IBM[®] Systems Director Management Console (SDMC), or the Integrated Virtualization Manager (IVM) to move an active or inactive logical partition from one server to another. You cannot migrate a mobile partition from a system that is managed by an HMC to a system that is managed by IVM. Similarly you cannot migrate a mobile partition from a system that is managed by IVM to a system that is managed by an HMC.

Because the HMC and SDMC always move the last activated profile, an inactive logical partition that has never been activated cannot be moved. For inactive partition mobility, you can either select the partition state defined in the hypervisor, or select the configuration data defined in the last activated profile on the source server. Use the IVM to move a logical partition that has never been activated.

You cannot perform Live Partition Mobility that is both bidirectional and concurrent. For example:

- When you are moving a mobile partition from the source server to the destination server, you cannot move another mobile partition from the destination server to the source server.
- When you are moving a mobile partition from the source server to the destination server, you cannot move another mobile partition from the destination server to some other server.

Related information:

[🔗 DB2 and System p virtualization: Performance and best practices](#)

[🔗 DB2 and the Live Partition Mobility feature of PowerVM on IBM System p using storage area network \(SAN\) storage](#)

[🔗 IBM PowerVM Live Partition Mobility](#)

What's new in Live Partition Mobility

Read about new or changed information in Live Partition Mobility since the previous update of this topic collection.

June 2016

- The following topic is new for N_Port ID Virtualization (NPIV) partition migration validation:
 - “NPIV LUN or disk level validation” on page 39
- The following topic is updated for NPIV partition migration validation:
 - “Specifying the attributes for a partition mobility operation by using the VIOS” on page 37

April 2014

- The following topics were updated for mobile partitions that have single root I/O virtualization (SR-IOV) logical ports:
 - “Configuration validation for partition mobility” on page 6
 - “HMC-managed systems: Preparing the source and destination servers for partition mobility” on page 48
 - “Preparing the mobile partition for partition mobility” on page 63
 - “Moving the mobile partition with HMC” on page 83

October 2013

The following updates have been made to the content.

- The following topic was updated for shared storage pools:
 - “Preparing the source and destination Virtual I/O Server logical partitions for partition mobility” on page 59
- The following topic was updated for virtual Ethernet adapters:
 - “Preparing the network configuration for partition mobility” on page 71
- The following topic was updated for physical volumes:
 - “Preparing the source and destination Virtual I/O Server logical partitions for partition mobility” on page 59
- The following topics were updated for Peer-to-Peer Remote Copy (PPRC) and N_Port ID Virtualization (NPIV adapters):
 - “Preparing the virtual SCSI configuration for partition mobility” on page 74
 - “Preparing the virtual Fiber Channel configuration for partition mobility” on page 79
- The following topic is new for improvement in partition mobility performance:
 - “Improving the performance of partition mobility” on page 56
 - “Server evacuation” on page 57
- The following topic was updated for synchronization of the current configuration capability:
 - “Configuration validation for partition mobility” on page 6

March 2013

The following updates have been made to the content.

- The following topics were updated for the IBM PowerLinux™ 7R1 (8246-L1D and 8246-L1T) and the IBM PowerLinux 7R2 (8246-L2D and 8246-L2T).
 - “HMC-managed systems: Preparing the source and destination servers for partition mobility” on page 48
 - “IVM-managed systems: Preparing the source and destination servers for partition mobility” on page 123
- The following topics were updated for mobile partitions that have SR-IOV logical ports:
 - “Configuration validation for partition mobility” on page 6
 - “HMC-managed systems: Preparing the source and destination servers for partition mobility” on page 48
- The following topics are new for mobile partitions that use the virtual server network (VSN):
 - “Verifying that the destination server supports the virtual server network” on page 127
 - “Determining the virtual Ethernet switch name and mode in the destination server” on page 128
 - “Configuring the Virtual I/O Server for the VSN capability” on page 132
- The following topics were updated for mobile partitions that use VSN:

- “Configuration validation for partition mobility” on page 6
- “HMC-managed systems: Preparing the source and destination servers for partition mobility” on page 48
- “Preparing the HMC for partition mobility” on page 58
- “Preparing the network configuration for partition mobility” on page 71
- “Resuming the suspended mobile partition with HMC” on page 86

February 2013

The following updates have been made to the content.

- “Configuration validation for partition mobility” on page 6

October 2012

The following updates have been made to the content.

- Added information about the IBM PowerLinux 7R1 (8246-L1C and 8246-L1D) servers.
 - “HMC-managed systems: Preparing the source and destination servers for partition mobility” on page 48
- The following information is new for mobile partitions that are configured with processing units to virtual processor ratio of less than 0.1 and greater than or equal to 0.05:
 - “Verifying the processor-level hardware capabilities of the destination server” on page 55
- The following information is updated for mobile partitions that are configured with processing units to virtual processor ratio of less than 0.1 and greater than or equal to 0.05:
 - “Configuration validation for partition mobility” on page 6
 - “HMC-managed systems: Preparing the source and destination servers for partition mobility” on page 48
 - “Preparing the HMC for partition mobility” on page 58
- The following information is new for the increased number of concurrent migration operations:
 - “Live Partition Mobility pseudodevice” on page 37
 - “Specifying the attributes for a partition mobility operation by using the VIOS” on page 37
 - “Specifying the attributes for a partition mobility operation by using the HMC” on page 38
 - “VIOS configuration options for partition mobility performance optimization” on page 40
- The following information is updated for the increased number of concurrent migration operations:
 - “Firmware support matrix for partition mobility” on page 51
- The following information was updated for secure IP tunnels:
 - “Configuring secure IP tunnels between the mover service partitions on the source and destination servers” on page 72
- The following information was updated for mobile partitions where you can specify the port name of the Fibre Channel during partition migration:
 - “Validating the configuration for partition mobility” on page 82
- The following information is updated for the Dynamic Platform Optimizer (DPO) feature:
 - “Configuration validation for partition mobility” on page 6
 - “Source and destination servers in a partition mobility environment” on page 29

July 2012

The following updates have been made to the content.

- Added information about the IBM PowerLinux 7R1 (8246-L1S) and IBM PowerLinux 7R2 (8246-L2S) servers.

- “HMC-managed systems: Preparing the source and destination servers for partition mobility” on page 48
- “IVM-managed systems: Preparing the source and destination servers for partition mobility” on page 123

Live Partition Mobility on HMC-managed systems

You can use the Hardware Management Console (HMC) to move an active or inactive logical partition from one server to another.

Partition mobility overview for HMC

You can learn about the benefits of partition mobility, how the Hardware Management Console (HMC) performs active and inactive partition mobility, and about the configuration that is required to successfully move a logical partition from one system to another.

Related tasks:

“Preparing for partition mobility” on page 48

You need to verify that the source and destination systems are configured correctly so that you can successfully move the mobile partition from the source system to the destination system. This includes verifying the configuration of the source and destination servers, the Hardware Management Console (HMC), the Virtual I/O Server logical partitions, the mobile partition, the virtual storage configuration, and the virtual network configuration.

Benefits of partition mobility

Partition mobility provides systems management flexibility and is designed to improve system availability.

For example:

- You can avoid planned outages for hardware or firmware maintenance by moving logical partitions to another server and then performing the maintenance. Partition mobility can help because you can use it to work around scheduled maintenance activities.
- You can avoid downtime for a server upgrade by moving logical partitions to another server and then performing the upgrade. This allows you to continue your work without disruption.
- If a server indicates a potential failure, you can move its logical partitions to another server before the failure occurs. Partition mobility can help avoid unplanned downtime.
- You can consolidate workloads running on several small, under used servers onto a single large server.
- You can move workloads from server to server to optimize resource use and workload performance within your computing environment. With active partition mobility, you can manage workloads with minimal downtime.
- For some systems, you can move applications from one server to an upgraded server by using PowerVM Live Partition Mobility, without affecting availability of the applications.

However, while partition mobility provides many benefits, it does not do the following functions:

- Partition mobility does not provide automatic workload balancing.
- Partition mobility does not provide a bridge to new functions. Logical partitions must be restarted and possibly reinstalled to take advantage of new features.

Partition mobility process

Learn about how the Hardware Management Console (HMC) moves an active or inactive logical partition from one server to another server.

The following table describes the steps that take place during the process of active and inactive partition mobility on the HMC.

Table 1. The steps involved in the process of active and inactive partition mobility on the HMC

Partition mobility step	Active mobility step	Inactive mobility step
1. You ensure that all requirements are satisfied and all preparation tasks are completed.	X	X
2. You shut down the mobile partition.		X
3. You initiate partition mobility by using the Partition Migration wizard on the HMC.	X	X
<p>4. The HMC extracts the physical device description for each physical adapter on the Virtual I/O Server logical partitions on the source server. The HMC uses the extracted information to determine whether the Virtual I/O Server (VIOS) partitions on the destination server can provide the mobile partition with the same virtual SCSI, virtual Ethernet, and virtual Fibre Channel configuration that exists on the source server. This operation includes verifying that the VIOS partitions on the destination server have enough available slots to accommodate the virtual adapter configuration of the mobile partition. The HMC uses all this information to generate a list of recommended virtual adapter mappings for the mobile partition on the destination server. Where possible, the HMC preserves the following configurations:</p> <ul style="list-style-type: none"> • Multipath I/O configurations. • Virtual slot assignments for virtual server adapters on the VIOS partitions. • User-defined names of the virtual target devices on the VIOS partitions. Partition mobility does not preserve vtscsix IDs. • User-defined adapter IDs for virtual server adapters on the VIOS partitions. <p>The HMC displays a list of recommended virtual adapter mappings (as well as all the possible virtual adapter mappings) for the mobile partition on the destination server. You can either use the virtual adapter mappings that are recommended by the HMC, or you can select different virtual adapter mappings for the mobile partition on the destination server.</p>	X	X
5. The HMC prepares the source and destination environments for partition mobility. This preparation includes using the virtual adapter mappings from step 4 for mapping the virtual adapters on the mobile partition to the virtual adapters on the VIOS partitions on the destination server.	X	X
6. The HMC transfers the logical partition state from the source environment to the destination environment. This transfer includes all the partition profiles that are associated with the mobile partition. The HMC modifies the active partition profile of the mobile partition to reflect the new virtual adapter mappings on the destination server.	<p>In active partition mobility, the following additional steps occur:</p> <ul style="list-style-type: none"> • The source mover service partition extracts the logical partition state information from the source server and sends it to the destination mover service partition over the network. • The destination mover service partition receives the logical partition state information and installs it on the destination server. 	X

Table 1. The steps involved in the process of active and inactive partition mobility on the HMC (continued)

Partition mobility step	Active mobility step	Inactive mobility step
7. The HMC suspends the mobile partition on the source server. The source mover service partition continues to transfer the logical partition state information to the destination mover service partition.	X	
8. The hypervisor resumes the mobile partition on the destination server.	X	
9. The HMC completes the migration. All resources that were consumed by the mobile partition on the source server are reclaimed by the source server: <ul style="list-style-type: none"> • The HMC removes the virtual SCSI adapters and the virtual Fibre Channel adapters (that were connected to the mobile partition) from the source VIOS partitions. • The HMC removes the virtual SCSI adapters, virtual Ethernet adapters, and virtual Fibre Channel adapters (that were connected to the mobile partition) from the partition profiles associated with the VIOS partitions on the source server. • For a mobile partition that uses shared memory, the HMC deactivates the paging space device that was used by the mobile partition and releases it so that it becomes available for other shared memory partitions to use. 	X	X
10. You activate the mobile partition on the destination server. (The processor and memory resources configured for the mobile partition remain unassigned until you activate the mobile partition on the destination server.)		X
11. You perform postrequisite tasks, such as adding dedicated I/O adapters to the mobile partition or adding the mobile partition to a partition workload group.	X	X

Configuration validation for partition mobility

You can learn about the tasks that the Partition Migration wizard on the Hardware Management Console (HMC) performs to validate your system configuration for active and inactive partition mobility.

Before you attempt to migrate an active logical partition, you must validate your environment. You can use the validation function on the HMC to validate your system configuration. If the HMC detects a configuration or connection problem, it displays an error message with information to help you resolve the problem.

The following tables list validation tasks that the HMC performs to verify that the source and destination systems are ready for active or inactive partition mobility.

General compatibility

Table 2. Validation tasks performed by the HMC to verify general compatibility for active and inactive partition mobility

Validation task	Active mobility task	Inactive mobility task
Checks that the HMC that manages the source server can successfully communicate with the HMC that manages the destination server, if they are different HMCs .	X	X

Table 2. Validation tasks performed by the HMC to verify general compatibility for active and inactive partition mobility (continued)

Validation task	Active mobility task	Inactive mobility task
Checks that the resource monitoring and control (RMC) connections are established.	Checks the RMC connections to the mobile partition, the source and destination Virtual I/O Server (VIOS) partitions, and the connection between the source and destination mover service partitions.	Checks the RMC connections to the source and destination VIOS partitions.
Checks mobility capability and compatibility.	Checks the source and destination servers, hypervisor, VIOS partitions, and mover service partitions.	Checks the VIOS and hypervisor.
Checks the number of current migrations against the number of supported migrations.	Checks the number of current active migrations against the number of supported active migrations.	Checks the number of current inactive migrations against the number of supported inactive migrations.

Server compatibility

Table 3. Validation tasks performed by the HMC to verify server compatibility for active and inactive partition mobility

Validation task	Active mobility task	Inactive mobility task
Checks that the necessary processing resources are available to create a shell logical partition on the destination system.	X	X
Checks that the necessary memory resources are available to create a shell logical partition on the destination system.	<ul style="list-style-type: none"> For a mobile partition that uses dedicated memory, checks that enough physical memory is available on the destination system. For a mobile partition that uses shared memory, checks that a shared memory pool is configured on the destination server and that it has enough physical memory to satisfy the entitled memory requirements of the mobile partition. 	For a mobile partition that uses dedicated memory, checks that enough physical memory is available on the destination system.

Table 3. Validation tasks performed by the HMC to verify server compatibility for active and inactive partition mobility (continued)

Validation task	Active mobility task	Inactive mobility task
<p>Checks that the necessary I/O adapter resources are available to create a shell logical partition on the destination system.</p> <p>During validation, the HMC extracts the device description for each virtual adapter on the VIOS partitions on the source server. The HMC uses the extracted information to determine whether the VIOS partitions on the destination server can provide the mobile partition with the same virtual SCSI, virtual Ethernet, and virtual Fibre Channel configuration that exists on the source server. This includes verifying that the VIOS partitions on the destination server have enough available slots to accommodate the virtual adapter configuration of the mobile partition.</p>	X	X
Checks that the logical memory block size is the same on the source and destination servers.	X	
If the mobile partition is capable of suspension, the HMC checks that the destination server supports partitions that are capable of suspension.	X	X
<p>When the firmware is at level 7.6, or later, you can configure virtual processors to use only 0.05 processing units per virtual processor. Consider the following restrictions when you migrate a partition to a server with the firmware at level 7.4, or earlier.</p> <p>Minimum processing units must be set to a value that results from the following calculation:</p> <p>0.1 × the minimum number of virtual processors that you select for the partition.</p> <p>Maximum processing units must be set to a value that results from the following calculation:</p> <p>0.1 × the maximum number of virtual processors that you select for the partition.</p> <p>Before you migrate partitions that use 0.05 processor units per virtual processor, you must ensure that the current ratio of assigned processor units to virtual processors is at least 0.1.</p>	X	X
If the mobile partition has single root I/O virtualization (SR-IOV) logical ports, that mobile partition cannot be migrated to the destination server. SR-IOV is a Peripheral Component Interconnect Special Interest Group specification to allow multiple partitions that are running simultaneously within a single computer to share a Peripheral Component Interconnect-Express (PCIe) device.	X	X

Table 3. Validation tasks performed by the HMC to verify server compatibility for active and inactive partition mobility (continued)

Validation task	Active mobility task	Inactive mobility task
As of HMC Version 7 Release 7.7.0, you can assign the Virtual Ethernet Port Aggregator (VEPA) switching mode to virtual Ethernet switches that are used by the virtual Ethernet adapters of the mobile partition. When the virtual Ethernet switch that is used by the virtual Ethernet adapter of the logical partition is enabled with the VEPA switching mode, then the logical partition uses virtual server network (VSN). If the mobile partition on the source server uses VSN, verify that the destination server also uses VSN.	X	X
When the HMC is at Version 7 Release 7.8.0, or later, the mobile partition supports synchronization of the current configuration capability. Verify that the HMC is at Version 7 Release 7.8.0, or later, on the destination server. For remote migration, if the HMC at the source server is at Version 7 Release 7.8.0, or later, and the HMC at the destination server is at a version earlier than Version 7 Release 7.8.0, then the current configuration profile is not visible on the destination server. If the HMC at the source server is at a version earlier than Version 7 Release 7.7.0, and the HMC at the destination server is at Version 7 Release 7.8.0, or later, then the current configuration profile is created on the destination server. When you connect a server to a HMC that is at a version earlier than Version 7 Release 7.8.0, after the server was earlier connected to a HMC at Version 7 Release 7.8.0, the last valid configuration profile is considered as a normal profile.	X	X

VIOS compatibility

Table 4. Validation tasks performed by the HMC to verify the source and destination VIOS partitions for active and inactive partition mobility

Validation task	Active mobility task	Inactive mobility task
Checks that all required I/O devices are connected to the mobile partition through a VIOS partition. That is, no physical adapters are assigned to the mobile partition and no virtual serial adapters are in virtual slots higher than 1.	X	X
Checks that no virtual SCSI disks are backed by logical volumes and that no virtual SCSI disks are attached to internal disks (not on the SAN).	X	X
Checks that the virtual SCSI disks assigned to the logical partition are accessible by the VIOS partitions on the destination server.		X
Checks that the reservation policies of the physical volumes are the same for the source and destination VIOS partitions.	X	X

Table 4. Validation tasks performed by the HMC to verify the source and destination VIOS partitions for active and inactive partition mobility (continued)

Validation task	Active mobility task	Inactive mobility task
Checks that the required virtual LAN IDs are available on the destination VIOS partitions can be preserved on the destination VIOS partitions.	X	X
Checks that the slot IDs of the virtual server adapters on the source VIOS partitions can be maintained on the destination VIOS partitions.	X	X
Checks that the user-defined names of the virtual target devices on the source VIOS partition can be maintained on the destination VIOS partition.	X	X
Checks that the user-defined adapter IDs of the virtual server adapters on the source VIOS partition can be maintained on the destination VIOS partition.	X	X
Checks that the redundancy configuration of the VIOS partitions on the source system can be maintained on the destination system. In some situations, you can move a logical partition to a destination system with less redundancy.	X	X
<p>For a mobile partition that uses shared memory, checks the following configuration:</p> <ul style="list-style-type: none"> • The number of active VIOS partitions (subsequently referred to as <i>paging VIOS partitions</i>) that are assigned to the shared memory pool on the destination server. • That an available paging space device exists on the destination server and that the device satisfies the following requirements: <ul style="list-style-type: none"> – It satisfies the redundancy preferences that you specify. – It meets the size requirements of the mobile partition (it is at least the size of the maximum logical memory of the mobile partition). <p>For example, you specify that the mobile partition uses redundant paging VIOS partitions on the destination server. You can move the mobile partition if the destination server provides the following configuration:</p> <ul style="list-style-type: none"> • Two paging VIOS partitions are assigned to the shared memory pool. • An available paging space device exists. • The paging space device meets the size requirements of the mobile partition. • Both paging VIOS partitions on the destination server have access to the paging space device. 	X	

Mobile partition compatibility

Table 5. Validation tasks performed by the HMC to verify that the mobile partition can successfully move to the destination server by using active or inactive partition mobility

Validation task	Active mobility task	Inactive mobility task
Checks that the operating system on the mobile partition is the Linux operating system.	X	X

Table 5. Validation tasks performed by the HMC to verify that the mobile partition can successfully move to the destination server by using active or inactive partition mobility (continued)

Validation task	Active mobility task	Inactive mobility task
Checks that the mobile partition has an active partition profile on the HMC.		X
Checks the mobile partition, its operating system, and its applications for migration capability.	X	
Checks that the mobile partition is not the redundant error path reporting logical partition.	X	X
Checks that the mobile partition is not in a partition workload group.	X	X
Checks the uniqueness of the virtual MAC addresses or the mobile partition.	X	X
Checks the state of the mobile partition.	Checks that the mobile partition state is Active or Running.	Checks that the mobile partition state is Not Activated.
Checks that the name of the mobile partition is not already in use on the destination server.	X	X
Checks that the mobile partition is not configured with barrier synchronization register (BSR) arrays.	X	
Checks that the mobile partition is not configured with huge pages.	X	
Checks that the mobile partition does not have a Host Ethernet Adapter (or Integrated Virtual Ethernet).	X	
Checks that the mobile partition is not performing a Dynamic Partition Optimizer (DPO) operation. DPO is a hypervisor function initiated by the HMC.	X	
Checks whether the mobile partition has any connected tape or optical devices as migration fails if any of these devices are connected.	X	X

Related tasks:

“Validating the configuration for partition mobility” on page 82

You can use the Partition Migration wizard on the Hardware Management Console (HMC) to validate the configuration of the source and destination systems for partition mobility. If the HMC detects a configuration or connection problem, it displays an error message with information to help you resolve the problem.

Related information:

 The Dynamic Platform Optimizer function

Logical partition attributes that change after the logical partition moves to the destination system

When you move a logical partition from one server to another, some of its attributes might change (such as the logical partition ID number) and some of its attributes remain the same (such as the logical partition configuration).

The following table describes the logical partition attributes that remain the same and the logical partition attributes that might change after you move a logical partition to the destination server.

Table 6. Logical partition attributes that might change or remain the same after a logical partition moves to the destination server

Attributes that remain the same	Attributes that might change
<ul style="list-style-type: none"> • The logical partition name • The logical partition type (dedicated processor or shared processor) • The logical partition configuration • The processor architecture • The Simultaneous Multi-Threading (SMT) state of each processor • The virtual MAC addresses, IP addresses, and LUN mapping to the target devices 	<ul style="list-style-type: none"> • The logical partition ID number • The machine type, model, and serial number • The model class of the underlying server • The processor version and type • The processor frequency • The affinity characteristics of the logical memory blocks (LMB) • The maximum number of hot pluggable and installed physical processors • The L1 and L2 cache size

Processor compatibility modes

Processor compatibility modes enable you to move logical partitions between servers that have different processor types without upgrading the operating environments installed in the logical partitions.

You can run several versions of the Linux and Virtual I/O Server operating environments in logical partitions on POWER5, POWER6®, POWER6+™, and POWER7® processor-based servers. Sometimes older versions of these operating environments do not support the capabilities that are available with new processors, thus limiting your flexibility to move logical partitions between servers that have different processor types.

A processor compatibility mode is a value assigned to a logical partition by the hypervisor that specifies the processor environment in which the logical partition can successfully operate. When you move a logical partition to a destination server that has a different processor type from the source server, the processor compatibility mode enables that logical partition to run in a processor environment on the destination server in which it can successfully operate. In other words, the processor compatibility mode enables the destination server to provide the logical partition with a subset of processor capabilities that are supported by the operating environment that is installed in the logical partition.

Related tasks:

“Verifying the processor compatibility mode of the mobile partition” on page 66

You can use the Hardware Management Console (HMC) to determine whether the processor compatibility mode of the mobile partition is supported on the destination server, and update the mode, if necessary, so that you can successfully move the mobile partition to the destination server.

“Verifying the processor compatibility mode of the mobile partition” on page 133

You can use the Integrated Virtualization Manager (IVM) to determine whether the processor compatibility mode of the mobile partition is supported on the destination server, and update the mode, if necessary, so that you can successfully move the mobile partition to the destination server.

Processor compatibility mode definitions:

You can learn about each processor compatibility mode and the servers on which each mode can run.

The following table describes each processor compatibility mode and the servers on which the logical partitions that use each processor compatibility mode can successfully operate.

Table 7. Processor compatibility modes

Processor compatibility mode	Description	Supported servers
POWER5	The POWER5 processor compatibility mode allows you to run operating-system versions that use all the standard features of the POWER5 processor.	Logical partitions that use the POWER5 processor compatibility mode can run on POWER5, POWER6, and POWER6+ processor-based servers. Restriction: A POWER6 processor cannot emulate all features of a POWER5 processor. Similarly, a POWER7 processor cannot emulate all features of a POWER6 or a POWER5 processor. For example, certain types of performance monitoring might not be available for a logical partition if the current processor compatibility mode of a logical partition is set to the POWER5 mode.
POWER6	The POWER6 processor compatibility mode allows you to run operating-system versions that use all the standard features of the POWER6 processor.	Logical partitions that use the POWER6 processor compatibility mode can run on POWER6, POWER6+, and POWER7 processor-based servers.
POWER6+	The POWER6+ processor compatibility mode allows you to run operating-system versions that use all the standard features of the POWER6+ processor.	Logical partitions that use the POWER6+ processor compatibility mode can run on POWER6+ and POWER7 processor-based servers.
POWER6 enhanced	The POWER6 enhanced processor compatibility mode allows you to run operating-system versions that use all the standard features of the POWER6 processor and also provides additional floating-point instructions to applications that use the POWER6 processor.	Logical partitions that use the POWER6 enhanced processor compatibility mode can run on POWER6 processor-based servers.
POWER6+ enhanced	The POWER6+ enhanced processor compatibility mode allows you to run operating-system versions that use all the standard features of the POWER6+ processor and also provides additional floating-point instructions to applications that use the POWER6+ processor.	Logical partitions that use the POWER6+ enhanced processor compatibility mode can run on POWER6+ processor-based servers.
POWER7	The POWER7 processor compatibility mode allows you to run operating-system versions that use all the standard features of the POWER7 processor.	Logical partitions that use the POWER7 processor compatibility mode can run on POWER7 processor-based servers.

Table 7. Processor compatibility modes (continued)

Processor compatibility mode	Description	Supported servers
default	The default processor compatibility mode is a preferred processor compatibility mode that enables the hypervisor to determine the current mode for the logical partition. When the preferred mode is set to default, the hypervisor sets the current mode to the most fully featured mode supported by the operating environment. In most cases, this is the processor type of the server on which the logical partition is activated. For example, assume that the preferred mode is set to default and the logical partition is running on a POWER7 processor-based server. The operating environment supports the POWER7 processor capabilities so the hypervisor sets the current processor compatibility mode to POWER7.	The servers on which logical partitions with the preferred processor compatibility mode of default can run depend on the current processor compatibility mode of the logical partition. For example, if the hypervisor determines that the current mode is POWER7, the logical partition can run on POWER7 processor-based servers.

Related concepts:

“Current and preferred processor compatibility modes”

The processor compatibility mode in which the logical partition currently operates is the *current* processor compatibility mode of the logical partition. The *preferred* processor compatibility mode of a logical partition is the mode in which you want the logical partition to operate.

“Enhanced processor compatibility modes” on page 17

The POWER6 enhanced and POWER6+ enhanced processor compatibility modes provide additional floating-point instructions to applications that use the POWER6 or POWER6+ processor.

“Examples: Using processor compatibility modes in partition mobility” on page 27

You can view examples of how processor compatibility modes are used when moving an active or inactive logical partition between servers with different processor types.

Related reference:

“Migration combinations of processor compatibility modes” on page 17

View all the combinations of the processor types of the source server, the processor types of the destination server, the current and preferred processor compatibility modes of the logical partition before the migration, and the current and preferred processor compatibility modes of the logical partition after the migration.

Current and preferred processor compatibility modes:

The processor compatibility mode in which the logical partition currently operates is the *current* processor compatibility mode of the logical partition. The *preferred* processor compatibility mode of a logical partition is the mode in which you want the logical partition to operate.

The hypervisor sets the current processor compatibility mode for a logical partition by using the following information:

- The processor features supported by the operating environment running in the logical partition.
- The preferred processor compatibility mode that you specify.

When you activate the logical partition, the hypervisor checks the preferred processor compatibility mode and determines whether the operating environment supports that mode. If the operating environment

supports the preferred processor compatibility mode, the hypervisor assigns the logical partition the preferred processor compatibility mode. If the operating environment does not support the preferred processor compatibility mode, the hypervisor assigns the logical partition the most fully featured processor compatibility mode that is supported by the operating environment.

The following table describes when each processor compatibility mode can be current mode or the preferred mode.

Table 8. Current and preferred processor compatibility modes

Processor compatibility mode	Can it be the current mode?	Can it be the preferred mode?
POWER5	Yes The POWER5 processor compatibility mode can be the current processor compatibility mode of a logical partition.	No You cannot specify POWER5 as the preferred processor compatibility mode. The only situation in which a logical partition will run in the POWER5 processor compatibility mode is when it is the only processor environment supported by the operating environment installed in the logical partition.
POWER6	Yes The POWER6 processor compatibility mode can be the current processor compatibility mode of a logical partition.	Yes You can specify POWER6 as the preferred processor compatibility mode for a logical partition.
POWER6+	Yes The POWER6+ processor compatibility mode can be the current processor compatibility mode of a logical partition.	Yes You can specify POWER6+ as the preferred processor compatibility mode for a logical partition.
POWER6 enhanced	Yes The POWER6 enhanced processor compatibility mode can be the current processor compatibility mode of a logical partition.	Yes You can specify POWER6 enhanced as the preferred processor compatibility mode for a logical partition.
POWER6+ enhanced	Yes The POWER6+ enhanced processor compatibility mode can be the current processor compatibility mode of a logical partition.	Yes You can specify POWER6+ enhanced as the preferred processor compatibility mode for a logical partition.
POWER7	Yes The POWER7 processor compatibility mode can be the current processor compatibility mode of a logical partition.	Yes You can specify POWER7 as the preferred processor compatibility mode for a logical partition.

Table 8. Current and preferred processor compatibility modes (continued)

Processor compatibility mode	Can it be the current mode?	Can it be the preferred mode?
default	No The default processor compatibility mode is a preferred processor compatibility mode.	Yes You can specify default as the preferred processor compatibility mode. Also, if you do not specify a preferred mode, the system automatically sets the preferred mode to default.

The following table shows the current and preferred processor compatibility modes supported on each server type.

Table 9. Processor compatibility modes supported by server type

Server processor type	Supported current modes	Supported preferred modes
POWER6+ processor-based server	POWER5, POWER6, POWER6+, POWER6+ enhanced	default, POWER6, POWER6+, POWER6+ enhanced
POWER6 processor-based server	POWER5, POWER6, POWER6 enhanced	default, POWER6, POWER6 enhanced
POWER7 processor-based server	POWER5, POWER6, POWER6+, POWER7	default, POWER6, POWER6+, POWER7

The preferred processor compatibility mode is the highest mode that the hypervisor can assign to a logical partition. If the operating environment installed in the logical partition does not support the preferred mode, the hypervisor can set the current mode to a lower mode than the preferred mode, but it cannot set the current mode to a higher mode than the preferred mode. For example, assume that a logical partition runs on a POWER7 processor-based server and you specify POWER7 as the preferred mode. The operating environment installed in the logical partition does not support the POWER7 processor capabilities, but it does support the POWER6 processor capabilities. When you activate the logical partition, the hypervisor assigns the POWER6 processor compatibility mode as the current mode for the logical partition because the POWER6 mode is the most fully featured mode that the operating environment supports and it is a lower mode than the preferred mode of POWER7.

You cannot dynamically change the current processor compatibility of a logical partition. To change the current processor compatibility mode, you must change the preferred processor compatibility mode, shut down the logical partition, and restart the logical partition. The hypervisor attempts to set the current processor compatibility mode to the preferred mode that you specified.

When you move an active logical partition between servers with different processor types, both the current and preferred processor compatibility modes of the logical partition must be supported by the destination server. When you move an inactive logical partition between servers with different processor types, only the preferred mode of the logical partition must be supported by the destination server.

If you specify the default mode as the preferred mode for an inactive logical partition, you can move that inactive logical partition to a server of any processor type. Because all servers support the default processor compatibility mode, you can move an inactive logical partition with the preferred mode of default to a server with any processor type. When the inactive logical partition is activated on the destination server, the preferred mode remains set to default, and the hypervisor determines the current mode for the logical partition.

Related concepts:

“Examples: Using processor compatibility modes in partition mobility” on page 27

You can view examples of how processor compatibility modes are used when moving an active or

inactive logical partition between servers with different processor types.

“Processor compatibility mode definitions” on page 12

You can learn about each processor compatibility mode and the servers on which each mode can run.

Related reference:

“Migration combinations of processor compatibility modes”

View all the combinations of the processor types of the source server, the processor types of the destination server, the current and preferred processor compatibility modes of the logical partition before the migration, and the current and preferred processor compatibility modes of the logical partition after the migration.

Enhanced processor compatibility modes:

The POWER6 enhanced and POWER6+ enhanced processor compatibility modes provide additional floating-point instructions to applications that use the POWER6 or POWER6+ processor.

Note: POWER7 processor-based servers do not support the enhanced mode.

If you want a logical partition to run in an enhanced mode, you must specify the enhanced mode as the preferred mode for the logical partition. If the operating environment supports the corresponding non-enhanced mode, then the hypervisor assigns the enhanced mode to the logical partition when you activate the logical partition. In other words, if you specify the POWER6+ enhanced mode as the preferred mode, and the operating environment supports the POWER6+ mode, the hypervisor assigns the POWER6+ enhanced mode to the logical partition when you activate the logical partition. Similarly, if you specify the POWER6 enhanced mode as the preferred mode, and the operating environment supports the POWER6 mode, the hypervisor assigns the POWER6 enhanced mode to the logical partition when you activate the logical partition.

Logical partitions in the POWER6 enhanced processor compatibility mode can only run on POWER6 processor-based servers, and logical partitions in the POWER6+ enhanced processor compatibility mode can only run on POWER6+ processor-based servers. Therefore, if a logical partition runs in the POWER6 enhanced mode, you can only move the logical partition to POWER6 processor-based servers. Likewise, if a logical partition runs in the POWER6+ enhanced mode, you can only move the logical partition to POWER6+ processor-based servers. If you want to move a logical partition in the POWER6 enhanced processor compatibility mode to a POWER6+ processor-based server, then you need to change the preferred mode to the default or POWER6 processor compatibility mode and restart the logical partition.

Related concepts:

“Examples: Using processor compatibility modes in partition mobility” on page 27

You can view examples of how processor compatibility modes are used when moving an active or inactive logical partition between servers with different processor types.

“Processor compatibility mode definitions” on page 12

You can learn about each processor compatibility mode and the servers on which each mode can run.

Related reference:

“Migration combinations of processor compatibility modes”

View all the combinations of the processor types of the source server, the processor types of the destination server, the current and preferred processor compatibility modes of the logical partition before the migration, and the current and preferred processor compatibility modes of the logical partition after the migration.

Migration combinations of processor compatibility modes:

View all the combinations of the processor types of the source server, the processor types of the destination server, the current and preferred processor compatibility modes of the logical partition before the migration, and the current and preferred processor compatibility modes of the logical partition after the migration.

Related concepts:

“Examples: Using processor compatibility modes in partition mobility” on page 27

You can view examples of how processor compatibility modes are used when moving an active or inactive logical partition between servers with different processor types.

“Enhanced processor compatibility modes” on page 17

The POWER6 enhanced and POWER6+ enhanced processor compatibility modes provide additional floating-point instructions to applications that use the POWER6 or POWER6+ processor.

“Current and preferred processor compatibility modes” on page 14

The processor compatibility mode in which the logical partition currently operates is the *current* processor compatibility mode of the logical partition. The *preferred* processor compatibility mode of a logical partition is the mode in which you want the logical partition to operate.

“Processor compatibility mode definitions” on page 12

You can learn about each processor compatibility mode and the servers on which each mode can run.

Migration combinations of processor compatibility modes for active partition mobility:

When you move an active logical partition between servers with different processor types, both the current and preferred processor compatibility modes of the logical partition must be supported by the destination server.

The following tables describe the processor compatibility mode combinations for active migrations. They show the processor type of the source server and the preferred and current processor compatibility modes of the logical partition on the source server before the migration. They also show the processor type of the destination server and the preferred and current processor compatibility modes of the logical partition on the destination server after the migration.

Table 10. Processor compatibility mode combinations for active migrations of POWER7 processor-based servers

Source environment			Destination environment		
Source server	Preferred mode before migration	Current mode before migration	Destination server	Preferred mode after migration	Current mode after migration
POWER7 processor-based server	default	POWER7, POWER6+, POWER6, or POWER5	POWER7 processor-based server	default	POWER7, POWER6+, POWER6, or POWER5
POWER7 processor-based server	POWER7	POWER7, POWER6+, POWER6, or POWER5	POWER7 processor-based server	POWER7	POWER7, POWER6+, POWER6, or POWER5
POWER7 processor-based server	POWER6+	POWER6+, POWER6, or POWER5	POWER7 processor-based server	POWER6+	POWER6+, POWER6, or POWER5
POWER7 processor-based server	POWER6	POWER6 or POWER5	POWER7 processor-based server	POWER6	POWER6 or POWER5

Table 10. Processor compatibility mode combinations for active migrations of POWER7 processor-based servers (continued)

Source environment			Destination environment		
Source server	Preferred mode before migration	Current mode before migration	Destination server	Preferred mode after migration	Current mode after migration
POWER7 processor-based server	default	POWER7, POWER6+, POWER6, or POWER5	POWER6+ processor-based server	default	If the current mode on the source server is POWER7, you cannot migrate the logical partition because the destination server does not support the current mode (POWER7). If the current mode on the source server is POWER6+, POWER6, or POWER5, then the current mode on the destination server is POWER6+, POWER6, or POWER5.
POWER7 processor-based server	POWER7	POWER7, POWER6+, POWER6, or POWER5	POWER6+ processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER7).	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER7).

Table 10. Processor compatibility mode combinations for active migrations of POWER7 processor-based servers (continued)

Source environment			Destination environment		
Source server	Preferred mode before migration	Current mode before migration	Destination server	Preferred mode after migration	Current mode after migration
POWER7 processor-based server	default	POWER7, POWER6+, POWER6, or POWER5	POWER6 processor-based server	default	If the current mode on the source server is POWER7 or POWER6+, you cannot migrate the logical partition because the destination server does not support the current mode (POWER7 or POWER6+). If the current mode on the source server is POWER6 or POWER5, then the current mode on the destination server is POWER6 or POWER5.
POWER7 processor-based server	POWER6+	POWER6+, POWER6, or POWER5	POWER6+ processor-based server	POWER6+	POWER6+, POWER6, or POWER5
POWER7 processor-based server	POWER6	POWER6 or POWER5	POWER6+ processor-based server	POWER6	POWER6 or POWER5
POWER7 processor-based server	POWER7 or POWER6+	POWER7, POWER6+, POWER6, or POWER5	POWER6 processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER7 or POWER6+).	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER7 or POWER6+).
POWER7 processor-based server	POWER6	POWER6 or POWER5	POWER6 processor-based server	POWER6	POWER6 or POWER5

Table 11. Processor compatibility mode combinations for active migrations of POWER6+ processor-based servers

Source environment			Destination environment		
Source server	Preferred mode before migration	Current mode before migration	Destination server	Preferred mode after migration	Current mode after migration
POWER6+ processor-based server	default	POWER6+, POWER6, or POWER5	POWER6+ processor-based server	default	POWER6+, POWER6, or POWER5

Table 11. Processor compatibility mode combinations for active migrations of POWER6+ processor-based servers (continued)

Source environment			Destination environment		
Source server	Preferred mode before migration	Current mode before migration	Destination server	Preferred mode after migration	Current mode after migration
POWER6+ processor-based server	POWER6+	POWER6+, POWER6, or POWER5	POWER6+ processor-based server	POWER6+	POWER6+, POWER6, or POWER5
POWER6+ processor-based server	POWER6+ enhanced	POWER6+ enhanced or POWER5	POWER6+ processor-based server	POWER6+ enhanced	POWER6+ enhanced or POWER5
POWER6+ processor-based server	POWER6	POWER6 or POWER5	POWER6+ processor-based server	POWER6	POWER6 or POWER5
POWER6+ processor-based server	default	POWER6+, POWER6, or POWER5	POWER6 processor-based server	default	If the current mode on the source server is POWER6+, you cannot migrate the logical partition because the destination server does not support the current mode (POWER6+). If the current mode on the source server is POWER6 or POWER5, then the current mode on the destination server is POWER6 or POWER5.
POWER6+ processor-based server	POWER6+	POWER6+, POWER6, or POWER5	POWER6 processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+).	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+).
POWER6+ processor-based server	POWER6+ enhanced	POWER6+ enhanced or POWER5	POWER6 processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+ enhanced).	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+ enhanced).

Table 11. Processor compatibility mode combinations for active migrations of POWER6+ processor-based servers (continued)

Source environment			Destination environment		
Source server	Preferred mode before migration	Current mode before migration	Destination server	Preferred mode after migration	Current mode after migration
POWER6+ processor-based server	POWER6	POWER6 or POWER5	POWER6 processor-based server	POWER6	POWER6 or POWER5
POWER6+ processor-based server	default	POWER6+, POWER6, or POWER5	POWER7 processor-based server	default	POWER7 (after you restart the logical partition), POWER6+, POWER6, or POWER5
POWER6+ processor-based server	POWER6+	POWER6+, POWER6, or POWER5	POWER7 processor-based server	POWER6+	POWER6+, POWER6, or POWER5
POWER6+ processor-based server	POWER6+ enhanced	POWER6+ enhanced or POWER5	POWER7 processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+ enhanced)	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+ enhanced)
POWER6+ processor-based server	POWER6	POWER6 or POWER5	POWER7 processor-based server	POWER6	POWER6 or POWER5

Table 12. Processor compatibility mode combinations for active migrations of POWER6 processor-based servers

Source environment			Destination environment		
Source server	Preferred mode before migration	Current mode before migration	Destination server	Preferred mode after migration	Current mode after migration
POWER6 processor-based server	default	POWER6 or POWER5	POWER6 processor-based server	default	POWER6 or POWER5
POWER6 processor-based server	POWER6	POWER6 or POWER5	POWER6 processor-based server	POWER6	POWER6 or POWER5
POWER6 processor-based server	POWER6 enhanced	POWER6 enhanced or POWER5	POWER6 processor-based server	POWER6 enhanced	POWER6 enhanced or POWER5
POWER6 processor-based server	default	POWER6 or POWER5	POWER6+ processor-based server	default	POWER6+ (after you restart the logical partition), POWER6, or POWER5
POWER6 processor-based server	POWER6	POWER6 or POWER5	POWER6+ processor-based server	POWER6	POWER6 or POWER5

Table 12. Processor compatibility mode combinations for active migrations of POWER6 processor-based servers (continued)

Source environment			Destination environment		
Source server	Preferred mode before migration	Current mode before migration	Destination server	Preferred mode after migration	Current mode after migration
POWER6 processor-based server	POWER6 enhanced	POWER6 enhanced or POWER5	POWER6+ processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6 enhanced).	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6 enhanced).
POWER6 processor-based server	default	POWER6 or POWER5	POWER7 processor-based server	default	POWER7 (after you restart the logical partition), POWER6 or POWER5
POWER6 processor-based server	POWER6	POWER6 or POWER5	POWER7 processor-based server	POWER6	POWER6 or POWER5
POWER6 processor-based server	POWER6 enhanced	POWER6 enhanced or POWER5	POWER7 processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6 enhanced)	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6 enhanced)

Related reference:

“Migration combinations of processor compatibility modes for inactive partition mobility”

When you move an inactive logical partition between servers with different processor types, only the preferred mode of the logical partition must be supported by the destination server.

“Migration combinations for version 1.5, and earlier, of the IVM” on page 114

Learn about the processor compatibility mode combinations for migrations where versions 1.5 (and earlier) of the Integrated Virtualization Manager (IVM) manage the source server and versions 2.1 (and later) of the IVM manage the destination server.

Migration combinations of processor compatibility modes for inactive partition mobility:

When you move an inactive logical partition between servers with different processor types, only the preferred mode of the logical partition must be supported by the destination server.

The following tables describe the processor compatibility mode combinations for inactive migrations. They show the processor type of the source server and the preferred processor compatibility modes of the logical partition on the source server before the migration. They also show the processor type of the destination server and the preferred and current processor compatibility modes of the logical partition on the destination server after the migration.

Table 13. Processor compatibility mode combinations for inactive migrations of POWER7 processor-based servers

Source environment		Destination environment		
Source server	Preferred mode before migration	Destination server	Preferred mode before migration	Current mode after migration
POWER7 processor-based server	default	POWER7 processor-based server	default	POWER7, POWER6+, POWER6, or POWER5
POWER7 processor-based server	POWER7	POWER7 processor-based server	POWER7	POWER7, POWER6+, POWER6, or POWER5
POWER7 processor-based server	POWER6+	POWER7 processor-based server	POWER6+	POWER6+, POWER6, or POWER5
POWER7 processor-based server	POWER6	POWER7 processor-based server	POWER6	POWER6 or POWER5
POWER7 processor-based server	default	POWER6+ processor-based server	default	POWER6+, POWER6, or POWER5
POWER7 processor-based server	POWER6+	POWER6+ processor-based server	POWER6+	POWER6+, POWER6, or POWER5
POWER7 processor-based server	POWER6	POWER6+ processor-based server	POWER6	POWER6 or POWER5
POWER7 processor-based server	POWER7	POWER6+ processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER7).	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER7).
POWER7 processor-based server	default	POWER6 processor-based server	default	POWER6 or POWER5
POWER7 processor-based server	POWER7 or POWER6+	POWER6 processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER7 or POWER6+).	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER7 or POWER6+).
POWER7 processor-based server	POWER6	POWER6 processor-based server	POWER6	POWER6 or POWER5

Table 14. Processor compatibility mode combinations for inactive migrations of POWER6+ processor-based servers

Source environment		Destination environment		
Source server	Preferred mode before migration	Destination server	Preferred mode before migration	Current mode after migration
POWER6+ processor-based server	default	POWER6+ processor-based server	default	POWER6+, POWER6, or POWER5
POWER6+ processor-based server	POWER6+	POWER6+ processor-based server	POWER6+	POWER6+, POWER6, or POWER5
POWER6+ processor-based server	POWER6	POWER6+ processor-based server	POWER6	POWER6 or POWER5
POWER6+ processor-based server	POWER6+ enhanced	POWER6+ processor-based server	POWER6+ enhanced	POWER6+ enhanced or POWER5
POWER6+ processor-based server	default	POWER6 processor-based server	default	POWER6 or POWER5
POWER6+ processor-based server	POWER6+	POWER6 processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+).	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+).
POWER6+ processor-based server	POWER6	POWER6 processor-based server	POWER6	POWER6 or POWER5
POWER6+ processor-based server	POWER6+ enhanced	POWER6 processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+ enhanced).	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+ enhanced).
POWER6+ processor-based server	default	POWER7 processor-based server	default	POWER7 (after you restart the logical partition), POWER6+, POWER6, or POWER5
POWER6+ processor-based server	POWER6+	POWER7 processor-based server	POWER6+	POWER6+, POWER6, or POWER5
POWER6+ processor-based server	POWER6+ enhanced	POWER7 processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+ enhanced)	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+ enhanced)

Table 14. Processor compatibility mode combinations for inactive migrations of POWER6+ processor-based servers (continued)

Source environment		Destination environment		
Source server	Preferred mode before migration	Destination server	Preferred mode before migration	Current mode after migration
POWER6+ processor-based server	POWER6	POWER7 processor-based server	POWER6	POWER6 or POWER5

Table 15. Processor compatibility mode combinations for inactive migrations of POWER6 processor-based servers

Source environment		Destination environment		
Source server	Preferred mode before migration	Destination server	Preferred mode before migration	Current mode after migration
POWER6 processor-based server	default	POWER6 processor-based server	default	POWER6 or POWER5
POWER6 processor-based server	POWER6	POWER6 processor-based server	POWER6	POWER6 or POWER5
POWER6 processor-based server	POWER6 enhanced	POWER6 processor-based server	POWER6 enhanced	POWER6 enhanced or POWER5
POWER6 processor-based server	default	POWER6+ processor-based server	default	POWER6+, POWER6, or POWER5
POWER6 processor-based server	POWER6	POWER6+ processor-based server	POWER6	POWER6 or POWER5
POWER6 processor-based server	POWER6 enhanced	POWER6+ processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6 enhanced).	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6 enhanced).
POWER6 processor-based server	default	POWER7 processor-based server	default	POWER7 (after you restart the logical partition), POWER6, or POWER5
POWER6 processor-based server	POWER6	POWER7 processor-based server	POWER6	POWER6 or POWER5
POWER6 processor-based server	POWER6 enhanced	POWER7 processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6 enhanced)	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6 enhanced)

Related reference:

“Migration combinations of processor compatibility modes for active partition mobility” on page 18
When you move an active logical partition between servers with different processor types, both the current and preferred processor compatibility modes of the logical partition must be supported by the destination server.

“Migration combinations for version 1.5, and earlier, of the IVM” on page 114

Learn about the processor compatibility mode combinations for migrations where versions 1.5 (and earlier) of the Integrated Virtualization Manager (IVM) manage the source server and versions 2.1 (and later) of the IVM manage the destination server.

Examples: Using processor compatibility modes in partition mobility:

You can view examples of how processor compatibility modes are used when moving an active or inactive logical partition between servers with different processor types.

Moving an active logical partition from a POWER6 processor-based server to a POWER7 processor-based server

You want to move an active logical partition from a POWER6 processor-based server to a POWER7 processor-based server so that the logical partition can take advantage of the additional capabilities available with the POWER7 processor.

To accomplish this task, perform the following steps:

1. Set the preferred processor compatibility mode to the default mode. When you activate the logical partition on the POWER6 processor-based server, it runs in the POWER6 mode.
2. Move the logical partition to the POWER7 processor-based server. Both the current and preferred modes remain unchanged for the logical partition until you restart the logical partition.
3. Restart the logical partition on the POWER7 processor-based server. The hypervisor evaluates the configuration. Because the preferred mode is set to default and the logical partition now runs on a POWER7 processor-based server, the highest mode available is the POWER7 mode. The hypervisor determines that the most fully featured mode supported by the operating environment installed in the logical partition is the POWER7 mode and changes the current mode of the logical partition to the POWER7 mode.

At this point, the current processor compatibility mode of the logical partition is the POWER7 mode and the logical partition runs on the POWER7 processor-based server.

Moving the active logical partition back to the POWER6 processor-based server

A problem arises and you need to move the active logical partition back to the POWER6 processor-based server. Because the logical partition now runs in the POWER7 mode and the POWER7 mode is not supported on the POWER6 processor-based server, you need to adjust the preferred mode for the logical partition so that the hypervisor can reset the current mode to a mode that is supported by the POWER6 processor-based server.

To move the logical partition back to the POWER6 processor-based server, perform the following steps:

1. Change the preferred mode from the default mode to the POWER6 mode.
2. Restart the logical partition on the POWER7 processor-based server. The hypervisor evaluates the configuration. Because the preferred mode is set to POWER6, the hypervisor does not set the current mode to a higher mode than POWER6. Remember, the hypervisor first determines whether it can set the current mode to the preferred mode. If not, it determines whether it can set the current mode to the next highest mode, and so on. In this case, the operating environment supports the POWER6 mode, so the hypervisor sets the current mode to the POWER6 mode.

3. Now that the logical partition runs in the POWER6 mode and the POWER6 mode is supported on the POWER6 processor-based server, move the logical partition back to the POWER6 processor-based server.

Moving an active logical partition between different processor types without making configuration changes

Depending on how often you need to move logical partitions, you might want to maintain the flexibility to move an active logical partition between a POWER6 processor-based server and a POWER7 processor-based server so that you can move the logical partition back and forth without making configuration changes. The easiest way to maintain this type of flexibility is to determine the processor compatibility mode supported on both the source and destination servers and set the preferred processor compatibility mode of the logical partition to the highest mode supported by both servers.

To accomplish this flexibility, perform the following steps:

1. Set the preferred processor compatibility mode to the POWER6 mode because the POWER6 mode is the highest mode supported by both POWER6 processor-based servers and POWER7 processor-based servers.
2. Move the logical partition from the POWER6 processor-based server to the POWER7 processor-based server.
3. Restart the logical partition on the POWER7 processor-based server. The hypervisor evaluates the configuration. Remember, the hypervisor does not set the current mode higher than the preferred mode. First, the hypervisor determines whether it can set the current mode to the preferred mode. If not, it then determines whether it can set the current mode to the next highest mode, and so on. In this case, the operating environment supports the POWER6 mode, so the hypervisor sets the current mode to the POWER6 mode.
4. Do not make any configuration changes to move the logical partition back to the POWER6 processor-based server because the POWER6 mode is supported on the POWER6 processor-based server.
5. Move the logical partition back to the POWER6 processor-based server.
6. Restart the logical partition on the POWER6 processor-based server. The hypervisor evaluates the configuration. The hypervisor determines that the operating environment supports the preferred mode of POWER6, and sets the current mode to the POWER6 mode.

Moving an inactive logical partition between servers with different processor types

The same logic from the previous examples applies to inactive partition mobility, except inactive partition mobility does not need the current processor compatibility mode of the logical partition because the logical partition is inactive. After you move an inactive logical partition to the destination server and activate that logical partition on the destination server, the hypervisor evaluates the configuration and sets the current mode for the logical partition just like it does when you restart a logical partition after active partition mobility. The hypervisor attempts to set the current mode to the preferred mode. If it cannot, it checks the next highest mode, and so on.

Related concepts:

“Enhanced processor compatibility modes” on page 17

The POWER6 enhanced and POWER6+ enhanced processor compatibility modes provide additional floating-point instructions to applications that use the POWER6 or POWER6+ processor.

“Current and preferred processor compatibility modes” on page 14

The processor compatibility mode in which the logical partition currently operates is the *current* processor compatibility mode of the logical partition. The *preferred* processor compatibility mode of a logical partition is the mode in which you want the logical partition to operate.

“Processor compatibility mode definitions” on page 12

You can learn about each processor compatibility mode and the servers on which each mode can run.

Related reference:

“Migration combinations of processor compatibility modes” on page 17

View all the combinations of the processor types of the source server, the processor types of the destination server, the current and preferred processor compatibility modes of the logical partition before the migration, and the current and preferred processor compatibility modes of the logical partition after the migration.

Partition mobility environment

You can learn about each component of the partition mobility environment and its contribution in enabling successful partition mobility. Components of the partition mobility environment include the source and destination servers, the Hardware Management Console (HMC), the source and destination Virtual I/O Server logical partitions, the mobile partition, the networking configuration, and the storage configuration.

Related tasks:

“Preparing for partition mobility” on page 48

You need to verify that the source and destination systems are configured correctly so that you can successfully move the mobile partition from the source system to the destination system. This includes verifying the configuration of the source and destination servers, the Hardware Management Console (HMC), the Virtual I/O Server logical partitions, the mobile partition, the virtual storage configuration, and the virtual network configuration.

Source and destination servers in a partition mobility environment:

Two servers are involved in partition mobility that is managed by a Hardware Management Console (HMC). The *source server* is the server from which you want to move the logical partition, and the *destination server* is the server to which you want to move the logical partition.

The source and destination servers must be POWER6 processor-based servers, or later, to participate in partition mobility. The destination server must have enough available processor and memory resources to allow the mobile partition to run on its server.

POWER7 processor-based servers with firmware at level 7.6, or later, can support the Dynamic Platform Optimizer (DPO) function. DPO is a hypervisor function initiated by the HMC. DPO rearranges the logical partition processors and memory on the system, to improve the affinity between processors and memory of the logical partition. When DPO is running, mobility operations targeting the system being optimized will be blocked. To continue with the migration, you must either wait for the DPO operation to complete, or manually stop the DPO operation.

Huge pages

Huge pages can improve performance in specific environments what require a high degree of parallelism, such as in DB2[®] partitioned database environments. You can specify the minimum, desired, and maximum number of huge pages to assign to a logical partition when you create the logical partition or partition profile.

A logical partition cannot participate in active partition mobility if huge pages are used. However, an inactive partition migration can be performed if the mobile partition uses huge pages. The partition profile will maintain the huge page resources, but the specified number of huge page resources may not be available on the destination server, in which case the logical partition will boot without some or all these huge pages after the inactive migration.

Barrier synchronization register

The barrier synchronization register (BSR) is a memory register that is located on certain processors based on POWER[®] technology.

A logical partition cannot participate in active partition migration if BSR is used. However, you can use inactive partition mobility if you do not want to disable BSR.

Shared memory pool

Shared memory is physical memory that is assigned to the shared memory pool and shared among multiple logical partitions. The *shared memory pool* is a defined collection of physical memory blocks that are managed as a single memory pool by the hypervisor. Logical partitions that you assign to the shared memory pool share the memory in the pool with other logical partitions that you assign to the pool.

If the mobile partition uses shared memory on the source server, the destination server must also have a shared memory pool to which the mobile partition can be assigned. If the mobile partition uses dedicated memory on the source server, it must also use dedicated memory on the destination server.

Inactive partition mobility policy


For inactive partition mobility, you can select one of the following configurations in the HMC for memory and processor-related settings of the mobile partition. If you are able to start the partition, and you select the current configuration as the mobility policy, then memory and processor-related settings are obtained from the partition state that is defined in the hypervisor. However, if you are unable to start the partition, or you select the last activated profile on the source server as the mobility policy, then memory and processor-related settings are obtained from the last activated profile on the source server. The mobility policy that you select applies to all inactive migrations, where the source server is the server on which you have set the policy.

For inactive partition mobility validation, the HMC either uses the hypervisor data or the last activated profile data to verify that the partition can be migrated to the destination server.

Related tasks:

“HMC-managed systems: Preparing the source and destination servers for partition mobility” on page 48
You need to verify that the source and destination servers are configured correctly so that you can successfully move the mobile partition from the source server to the destination server by using the Hardware Management Console (HMC). This includes tasks such as verifying the logical memory block size of the source and destination servers, and verifying the available memory and processor resources of the destination server.

Related information:

 [Overview of shared memory](#)

Hardware Management Console in a partition mobility environment:

Learn about the Hardware Management Console (HMC) and how you can use its Partition Migration wizard to move an active or inactive logical partition from one server to another server.

The HMC is a system that controls managed systems, including the management of logical partitions and the use of Capacity on Demand. Using service applications, the HMC communicates with managed systems to detect, consolidate, and send information to IBM for analysis.

Partition mobility can include one or more HMC as follows:

- Both the source and destination servers are managed by the same HMC (or redundant HMC pair). In this case, the HMC must be at version 7, release 7.1, or later.
- The source server is managed by one HMC and the destination server is managed by a different HMC. In this case, both the source HMC and the destination HMC must meet the following requirements:
 - The source HMC and the destination HMC must be connected to the same network so that they can communicate with each other.
 - The source HMC and the destination HMC must be at version 7, release 7.1, or later.

The HMC can handle multiple migrations simultaneously. However, the maximum number of concurrent partition migrations is limited by the processing capacity of the HMC.

The partition mobility wizard that is provided on the HMC helps you validate and complete a partition migration. The HMC determines the appropriate type of migration to use based on the state of the logical partition. If the logical partition is in the Running state, then the migration is active. If the logical partition is in the Not Activated state, then the migration is inactive. Before the migration starts, the HMC validates your logical partition environment. During this validation, the HMC determines if the migration will be successful. If the validation fails, the HMC provides error messages and suggestions to help you resolve the configuration problems.

Related tasks:

“Preparing the HMC for partition mobility” on page 58

You need to verify that the Hardware Management Console (HMC) that manage the source and destination servers are configured correctly so that you can move the mobile partition from the source server to the destination server.

Source and destination Virtual I/O Server logical partitions in a partition mobility environment:

Partition mobility that is managed by a Hardware Management Console (HMC) requires at least one Virtual I/O Server (VIOS) logical partition on the source server and at least one VIOS logical partition on the destination server.

Server partition

The mobile partition must receive storage and networking resources from the following sources:

- At least one VIOS logical partition on the source server.
- At least one VIOS logical partition on the destination server.

The VIOS logical partitions provide the mobile partition with access to the same storage from both the source and destination servers.

The mobile partition can access its physical storage through redundant VIOS logical partitions, a VIOS logical partition with redundant physical adapters, or both. In most cases, you must maintain the redundancy configuration of the VIOS logical partitions on the destination system. However, in some situations, you can move a logical partition to a destination system with less redundancy.

Where possible, partition mobility preserves the following configuration attributes:

- Slot IDs of the virtual server adapters
- User-defined names of virtual target devices
- User-defined adapter IDs of the virtual server adapters

Mover service partition

For active partition mobility, the following logical partitions must be designated as mover service partitions:

- At least one VIOS logical partition on the source server.
- At least one VIOS logical partition on the destination server

A *mover service partition* is a VIOS logical partition with the following characteristics:

- The mover service partition attribute indicates that the VIOS logical partition can support active partition migration.
- Both VIOS partitions must be at version 1.5 or later.

The source and destination mover service partitions communicate with each other over the network. On both the source and destination servers, the Virtual Asynchronous Services Interface (VASI) device provides communication between the mover service partition and the hypervisor. These connections facilitate active partition mobility as follows:

- On the source server, the mover service partition extracts the logical partition state information of the mobile partition from the hypervisor.
- The mover service partition on the source server sends the logical partition state information to the mover service partition on the destination server.
- On the destination server, the mover service partition installs the logical partition state information to the hypervisor.

Paging VIOS partition

A VIOS logical partition that is assigned to the shared memory pool (hereafter referred to as a *paging VIOS partition*) provides access to the paging space devices for the logical partitions that use shared memory.

You are not required to maintain the same number of paging VIOS partitions for the mobile partition from the source server to the destination server. For example, a mobile partition that uses redundant paging VIOS partitions on the source server can move to a destination server with only one paging VIOS partition assigned to the shared memory pool. Similarly, a mobile partition that uses a single paging VIOS partition on the source server can use redundant paging VIOS partitions on the destination server, if two paging VIOS partitions are assigned to the shared memory pool on the destination server. The following table describes these redundancy options in more detail.

When you validate the configuration for active partition mobility, the HMC checks that the paging VIOS partitions on the destination system have access to a paging space device that meets the size requirements of the mobile partition as well as the redundancy preferences that you specify. The HMC selects and assigns paging space devices to the mobile partition on the destination system by using the same process as used during partition activation. For details, see *Paging space devices on systems that are managed by an HMC*.

Table 16. Redundancy options for the paging VIOS partitions that are assigned to the mobile partition

Number of paging VIOS partitions that are used by the mobile partition on the source server	Number of paging VIOS partitions that are assigned to the shared memory pool on the destination server
<p>1</p> <p>The mobile partition uses a single paging VIOS partition to access its paging space device on the source system.</p>	<p>1</p> <p>Because there is only one paging VIOS partition that is assigned to the shared memory pool on the destination system, the mobile partition must continue to use a single paging VIOS partition to access a paging space device on the destination system.</p> <p>To successfully move the mobile partition in this situation, you can perform one of the following actions:</p> <ul style="list-style-type: none"> • Do not specify a redundancy preference. By default, the HMC attempts to maintain the current redundancy configuration on the destination system. In this case, the mobile partition continues to use a single paging VIOS partition to access a paging space device on the destination system. • Specify that the mobile partition does not use redundant paging VIOS partitions. The mobile partition continues to use a single paging VIOS partition to access a paging space device on the destination system. • Specify that the mobile partition uses redundant paging VIOS partitions, if possible. Use this option if you do not know whether the mobile partition can use redundant paging VIOS partitions on the destination system. The HMC examines the destination system to determine whether it is configured to support redundant paging VIOS partitions. In this case, the HMC finds that the mobile partition cannot use redundant paging VIOS partitions because only one paging VIOS partition is assigned to the shared memory pool on the destination server. Instead, the mobile partition continues to use a single paging VIOS partition to access a paging space device on the destination system.

Table 16. Redundancy options for the paging VIOS partitions that are assigned to the mobile partition (continued)

Number of paging VIOS partitions that are used by the mobile partition on the source server	Number of paging VIOS partitions that are assigned to the shared memory pool on the destination server
<p>1</p> <p>The mobile partition uses a single paging VIOS partition to access its paging space device on the source system.</p>	<p>2</p> <p>To successfully move the mobile partition in this situation, you can perform one of the following actions:</p> <ul style="list-style-type: none"> • Do not specify a redundancy preference. By default, the HMC attempts to maintain the current redundancy configuration on the destination system. In this case, the mobile partition continues to use a single paging VIOS partition to access a paging space device on the destination system. • Specify that the mobile partition does not use redundant paging VIOS partitions. The mobile partition continues to use a single paging VIOS partition to access a paging space device on the destination system. • Specify that the mobile partition uses redundant paging VIOS partitions, if possible. Use this option if you want the mobile partition to use redundant paging VIOS partitions on the destination system or if you do not know whether the mobile partition can use redundant paging VIOS partitions on the destination system. The HMC examines the destination system to determine whether it is configured to support redundant paging VIOS partitions. In this case, the HMC finds that the mobile partition can use redundant paging VIOS partitions because two paging VIOS partitions are assigned to the shared memory pool on the destination server. The mobile partition uses redundant paging VIOS partitions to access a paging space device on the destination system.

Table 16. Redundancy options for the paging VIOS partitions that are assigned to the mobile partition (continued)

Number of paging VIOS partitions that are used by the mobile partition on the source server	Number of paging VIOS partitions that are assigned to the shared memory pool on the destination server
<p>2</p> <p>The mobile partition uses redundant paging VIOS partitions to access its paging space device on the source system.</p>	<p>1</p> <p>Because only one paging VIOS partition is assigned to the shared memory pool on the destination server, the mobile partition cannot continue to use redundant paging VIOS partitions to access a paging space device on the destination system. Instead, it must use a single paging VIOS partition to access a paging space device.</p> <p>To successfully move the mobile partition in this situation, you can perform one of the following actions:</p> <ul style="list-style-type: none"> • Specify that the mobile partition does not use redundant paging VIOS partitions. The mobile partition uses a single paging VIOS partition to access a paging space device on the destination system. • Specify that the mobile partition uses redundant paging VIOS partitions, if possible. Use this option if you do not know whether the mobile partition can use redundant paging VIOS partitions on the destination system. The HMC examines the destination system to determine whether it is configured to support redundant paging VIOS partitions. In this case, the HMC finds that the mobile partition cannot use redundant paging VIOS partitions because only one paging VIOS partition is assigned to the shared memory pool on the destination server. Instead, the mobile partition uses a single paging VIOS partition to access a paging space device on the destination system.

Table 16. Redundancy options for the paging VIOS partitions that are assigned to the mobile partition (continued)

Number of paging VIOS partitions that are used by the mobile partition on the source server	Number of paging VIOS partitions that are assigned to the shared memory pool on the destination server
<p>2</p> <p>The mobile partition uses redundant paging VIOS partitions to access its paging space device on the source system.</p>	<p>2</p> <p>To successfully move the mobile partition in this situation, you can perform one of the following actions:</p> <ul style="list-style-type: none"> • Do not specify a redundancy preference. By default, the HMC attempts to maintain the current redundancy configuration on the destination system. In this case, the mobile partition continues to use redundant paging VIOS partitions to access a paging space device on the destination system. • Specify that the mobile partition does not use redundant paging VIOS partitions. The mobile partition uses a single paging VIOS partition to access a paging space device on the destination system. • Specify that the mobile partition uses redundant paging VIOS partitions, if possible. Use this option if you want the mobile partition to use redundant paging VIOS partitions on the destination system or if you do not know whether the mobile partition can use redundant paging VIOS partitions on the destination system. The HMC examines the destination system to determine whether it is configured to support redundant paging VIOS partitions. In this case, the HMC finds that the mobile partition can use redundant paging VIOS partitions because two paging VIOS partitions are assigned to the shared memory pool on the destination server. The mobile partition continues to use redundant paging VIOS partitions to access a paging space device on the destination system.

Related concepts:

“Network configuration in a partition mobility environment” on page 42

In partition mobility that is managed by the Hardware Management Console (HMC), the network between the source and destination servers is used to pass the mobile partition state information and other configuration data from the source environment to the destination environment. The mobile partition uses the virtual LAN for network access.

“Storage configuration in a partition mobility environment” on page 43

Learn about the virtual SCSI and virtual Fibre Channel configuration required for partition mobility that is managed by the Hardware Management Console (HMC).

Related tasks:

“Preparing the source and destination Virtual I/O Server logical partitions for partition mobility” on page 59

You must verify that the source and destination Virtual I/O Server (VIOS) logical partitions are configured correctly so that you can successfully move the mobile partition from the source server to the destination server by using the Hardware Management Console (HMC). This verification includes tasks such as verifying the version of the VIOS partitions and enabling the mover service partitions.

“Verifying that the destination shared memory pool contains an available paging space device” on page 62

You can verify that the shared memory pool on the destination server contains a paging space device that satisfies the size requirements and redundancy configuration of the mobile partition by using the

Hardware Management Console (HMC).

Related information:

 [Paging VIOS partition](#)

Live Partition Mobility pseudodevice:

The **vioslpm0** pseudodevice is created by default when you install Virtual I/O Server (VIOS) Version 2.2.2.0. You can use the attributes of the partition mobility pseudodevice to control active partition mobility operations. The pseudodevice saves the attributes that affect partition mobility operations.

Specifying the attributes for a partition mobility operation by using the VIOS:

You can specify the attributes for a partition mobility operation by using the Virtual I/O Server (VIOS). The attributes that are specified are saved in the **vioslpm0** pseudodevice.

The following list describes how to specify the attributes for the **vioslpm0** pseudodevice by using the VIOS command line.

You can list the attributes associated with the **vioslpm0** pseudodevice by running the following command, where *vioslpm0* is the name of the pseudodevice:

```
lsdev -dev vioslpm0 -attr
```

You can set the following attributes:

- The **cfg_msp_lpm_ops** attribute is used to control the maximum number of concurrent partition mobility operations that the VIOS can support. You can limit the number of concurrent partition mobility operations that the VIOS will run, based on the configuration and work load of the VIOS. For example, if the VIOS is configured with a single 1 GB network adapter, the value of the **cfg_msp_lpm_ops** attribute must be a value of 4. The default value for this attribute is 8 for VIOS version 2.2.2.0, or later; therefore, the VIOS version 2.2.2.0 supports up to eight concurrent partition mobility operations. To run the maximum number of supported partition mobility operations on the VIOS, this value must be set to the supported maximum number. The attribute value range is 1 - 8 for VIOS version 2.2.2.0, or later
- The **concurrency_lvl** attribute controls the amount of resources allocated for each partition mobility operation. The attribute value range is 1 - 5. For optimal performance, use a value of 1. When you have limited resources, use a value of 5. The default value is 3. It is suggested that the default value of 3 be used in all situations. If migrations fail because of memory limitations, changing the **concurrency_lvl** attribute to a higher value of 4 or 5 might resolve the problem, because the migration will use fewer memory resources in the VIOS partitions.
- The **lpm_msnap_succ** attribute indicates whether partition mobility trace data must be saved for migrations that complete successfully. This information is required by IBM support teams to analyze partition mobility performance problems. The default value is 1, which means that data from successful partition mobility operations is saved.
- The **tcp_port_high** and **tcp_port_low** attributes are used to control the range of ports you can select for partition mobility operations. By default, both attributes are set to zero, indicating that any of the 32,768 ephemeral ports on the VIOS can be used for partition mobility operation. When you set the port range, it is suggested that you allocate enough ports for the maximum number of concurrent partition mobility operations in addition to a few more. This helps prevent partition mobility operations from failing in the event that one or more of the ports are in use by other parts of the system. Two ports are used for each partition mobility operation.
- The **src_lun_val** attribute is used to enable and disable LUN level validation of N_Port ID Virtualization (NPIV) devices. This attribute had two possible values, *on* and *off*. When the attribute is set to *off*, LUN level validation is not performed, and when the attribute is set to *on*, LUN level validation is performed. For more information about disk level validation, see “NPIV LUN or disk level validation” on page 39.

- The **dest_lun_val** attribute is used to disable LUN level validation of NPIV devices for different operations and is relevant only when **src_lun_val** has the value *on* in the source VIOS. This attribute affects only the destination VIOS that is hosting the NPIV storage for remote restart and partition mobility operations. There are four allowed values for this attribute, *on*, *off*, *restart_off*, and *lpm_off*. By default the attribute is set to *restart_off*. This value disables LUN level validation for remote restart but allows it for partition mobility operations. Setting the attribute to *lpm_off* allows LUN level validation for remote restart operations but disables it for partition mobility operations. A value *on* allows LUN level validation for both partition mobility and remote restart and a value *off* disables LUN level validation for all operations. For more information on disk level validation, see “NPIV LUN or disk level validation” on page 39.
- The **max_val_cmds** attribute controls the number of command elements that are allocated for NPIV disk level validation. Higher values reduce the amount of time that is required to perform disk level validation, but they also allocate more resources and use more SAN bandwidth per physical port. It is recommended that the default value is used unless the user has more than 100 disks and the validation time is unacceptable since there is no performance advantage to changing this attribute if the client does not have more than 100 devices visible through the port. For more information on disk level validation, see “NPIV LUN or disk level validation” on page 39.

Table 17. Pseudodevice attributes and definition

Attribute	Value	Description	User modifiable
cfg_msp_lpm_ops	8	Number of concurrent partition mobility operations for the mover service partition	True
concurrency_lvl	3	Concurrency level	True
lpm_msnap_succ	1	Create a mini-snap (when a migration ends, the set of information related to a specific migration, that is gathered and packed on each mover service partition involved in the migration), for successful migrations	True
max_lpm_vasi	1	Maximum number of Virtual Asynchronous Services Interface (VASI) adapters used for partition mobility operations	False
max_vasi_ops	8	Maximum number of concurrent partition mobility operations per VASI	False
tcp_port_high	0	TCP highest ephemeral port	True
tcp_port_low	0	TCP lowest ephemeral port	True
src_lun_val	<i>off</i>	Enable or disable NPIV disk validation for remote restart	True
dest_lun_val	<i>restart_off</i>	Enable or disable NPIV disk validation for partition mobility	True
max_val_cmds	100	Change the number of commands that are allocated for NPIV LPM disk validation	True

As shown in the previous table, you can change the values of the attributes that are user modifiable. For example, to specify a value of 5 for the **cfg_msp_lpm_ops** attribute, run the following command:

```
chdev -dev vioslpm0 -attr cfg_msp_lpm_ops=5
```

Specifying the attributes for a partition mobility operation by using the HMC:

You can specify the attributes for a partition mobility operation by using the Hardware Management Console (HMC).

To specify the attributes for a partition mobility operation by using the HMC command line, complete the following steps:

1. To list the attributes associated with the partition mobility operation, run the following command:

where:

- *srcCecName* is the name of the server from which you want to move the mobile partition.
- *dstCecName* is the name of the server to which you want to move the mobile partition.
- *lparName* is the name of the logical partition to be migrated.

```
lslparmigr -r msp -m <srcCecName> -t <dstCecName> --filter "lpar_names=<lparName>"
```

2. Run the following command to modify the attributes of a partition mobility operation

```
migr1par -o set -r lpar -m <CecName> -p <lparName> -i "...."
```

You can modify the following attributes by using the **migr1par** command:

- **num_active_migrations_configured**
- **concurr_migration_perf_level**

For example:

- To set the number of concurrent active migrations that can be run to a value of 8, run the following command:

```
migr1par -o set -r lpar -m <CecName> -p <lparName> -i "num_active_migrations_configured=8"
```

The default value for this attribute is 4. To run the maximum number of supported partition mobility operations on the Virtual I/O Server (VIOS), set this value to the maximum number that is supported.

- To set the amount of resources allocated for each mobility operation to a value of 2, run the following command:

```
migr1par -o set -r lpar -m <CecName> -p <lparName> -i "concurr_migration_perf_level=2"
```

The range of the attribute value is 1 - 5. A value of 1 indicates optimal performance, and a value of 5 indicates limited resources. The default value is 3.

NPIV LUN or disk level validation:

With the Virtual I/O Server (VIOS) version 2.2.4.0 or earlier, partition mobility validation for N_Port ID Virtualization (NPIV) devices are checked only up to the port level. This resulted in the possibility of client failures if the actual disk mapped to the client on the source system were not properly mapped on the destination system. With the VIOS version 2.2.4.0, you can validate up to the disk mapping. To perform disk level validation, both the source and destination VIOS must be at level 2.2.4.0 or later, and the Hardware Management Console (HMC) must be at least at Version 7.4.4.

Disk validation can add a considerable amount of time to partition mobility validation for clients that are using NPIV disks. The amount of time that is required to validate NPIV devices up to the disk level depends on the number of disks that are mapped to a client. For larger configurations, the additional time that is spent in validation might have a noticeable impact on the overall time that is required to migrate the partition. Therefore, it is suggested that you might want to consider performing periodic partition mobility validation with LUN level validation enabled. Also, it would be prudent to plan the validation outside of scheduled maintenance windows and either skip validation, or run validation with LUN level validation disabled when partition mobility operations must be completed in a short period.

To enable disk level validation, the **src_lun_val** attributes in the Live Partition Mobility pseudodevice of the VIOS that is hosting the NPIV storage on the source system must be set to a value of *on* and the **dest_lun_val** attribute on the VIOS partitions that are hosting the NPIV storage on the destination system cannot be set to *lpm_off* or *off*.

Note:

- Since disk validation sends additional commands to the SAN, any instability in the SAN might result in validation failures where port level validation might have succeeded.
- Disk mapping validation is done during partition mobility validation, and is not done during migration. The migration phase of a partition mobility operation validates only up to the port level.
- When you are using the HMC graphical user interface, validation is always done for each partition mobility operation. You must keep this in mind before you enable disk level validation, particularly if the client has many disks.
- When the HMC command line interface is used, validation is performed only if the `-o` flag is set to the character `v` and migration is done only if the `—o` flag is set to the character `m`. These flags are mutually exclusive.

VIOS configuration options for partition mobility performance optimization:

Partition mobility operations require a significant amount of system resources to achieve maximum performance while maintaining client stability. Configure the source and destination mover server partitions to have a similar amount of processing capabilities, because the overall performance of the migration is limited by the mover server partition that is configured with fewer processing capabilities.

Configuring the VIOS firewall for partition mobility:

You must configure the Virtual I/O Server (VIOS) firewall to allow partition mobility before you enable the VIOS firewall.

Partition mobility operations fail when the VIOS firewall is enabled with default settings. The partition mobility operations fail because the firewall blocks the Internet Control Message Protocol (ICMP) that is required during partition mobility validation. Also, partition mobility fails because the firewall blocks ephemeral ports that are required for partition mobility. You must manually configure the VIOS firewall to prevent partition mobility failure.

To Add ICMP roles to the firewall configuration on all the Virtual I/O Servers, complete the following steps

1. Run the **oem_setup_env** command from the VIOS command line. Running this command provides a new environment to run other commands.
2. From the new environment, run the following commands:
 - a.

```
/usr/sbin/genfilt -v 4 -n 16 -a P
-s 0.0.0.0 -m 0.0.0.0 -d 0.0.0.0
-M 0.0.0.0 -g n -c icmp -o eq -p 0
-o any -P 0 -r L -w I -l N -t 0
-i all -D echo_reply
```
 - b.

```
/usr/sbin/genfilt -v 4 -n 16 -a P
-s 0.0.0.0 -m 0.0.0.0 -d 0.0.0.0
-M 0.0.0.0 -g n -c icmp -o eq -p 8
-o any -P 0 -r L -w I -l N -t 0
-i all -D echo_request
```
 - c. Run the **exit** command to return to the VIOS command line.
3. Reduce the range of ephemeral ports and create a role for each of the ephemeral ports in the firewall configuration.

For example, to reduce the range of ephemeral ports to nine, run the following commands from the VIOS command line:

```
chdev -dev vioslpm0 -attr tcp_port_high=40010
chdev -dev vioslpm0 -attr tcp_port_low=40001
```

Note: Live Partition Mobility uses two ephemeral ports per migration. The ephemeral port range is 0 - 32 KB and the network stack randomly selects the ports to be used for partition mobility operations.

With VIOS version 2.2.2.0, or later, the **tcp_port_high** and **tcp_port_low** attributes are used to control the range of ports that you can select for partition mobility operations. You can change the value by using the **chdev** command. You must use a range that is large enough to run the maximum number of concurrent partition mobility operations, and some more ports in case any of the ports are used by another program.

4. Enable the ports to be used in the VIOS firewall.

For example, to enable the ports one and two in the VIOS firewall, run the following commands from the VIOS command line:

```
viosecure -firewall allow -port 40001
viosecure -firewall allow -port 40002
```

Mobile partition managed by an HMC in a partition mobility environment:

A *mobile partition* is a logical partition that you want to move from the source server to the destination server. You can move a running mobile partition, or active mobile partition, or you can move a powered off mobile partition, or inactive mobile partition, from the source server to the destination server.

The HMC creates a migration profile for the mobile partition on the destination server that matches the current configuration of the logical partition. During the migration, the HMC migrates all the profiles associated with the mobile partition to the destination server. Only the current partition profile (or a new one, if specified) is converted during the migration process. This conversion includes mapping the client virtual SCSI slot and the client virtual Fibre Channel slot to the corresponding target virtual SCSI slot and the corresponding target virtual Fibre Channel slot on the destination Virtual I/O Server logical partitions, if required.

A logical partition cannot be migrated if any logical partition exists on the destination server with the same name. The HMC creates a migration profile containing the current state of the logical partition if you do not specify a profile name. The profile replaces the existing profile that was last used to activate the logical partition. If you specify an existing profile name, the HMC replaces that profile with the new migration profile. If you want to keep the existing profiles of the logical partition, specify a new and unique profile name before the migration begins.

For inactive partition mobility, the HMC provides you an option to select one of the following configurations for memory and processor-related settings of the mobile partition. If you are able to start the partition, and you select the current configuration as the mobility policy, then memory and processor-related settings are obtained from the partition state defined in the hypervisor. However, if you are unable to start the partition, or you select the last activated profile on the source server as the mobility policy, then memory and processor-related settings are obtained from the last activated profile on the source server. The mobility policy that you select applies to all inactive migrations where the source server is the server on which you have set the policy.

Considerations for configuring I/O

Do not assign any physical or required I/O adapters to a mobile partition using the active partition migration. All the I/O adapters on the mobile partition must be virtual devices. To remove the physical adapters on the mobile partition, you can use the dynamic logical partition removal task.

A mobile partition with dedicated adapters can participate in inactive partition mobility; however, the dedicated adapters will be removed from the partition profile. Thus, the logical partition will boot with only virtual I/O resources after an inactive migration. If dedicated I/O resources were assigned to the logical partition on the source server, these resources will become available when the logical partition is deleted from the source server.

Related tasks:

“Preparing the mobile partition for partition mobility” on page 63

You need to verify that the mobile partition is configured correctly so that you can successfully move it from the source server to the destination server by using the Hardware Management Console (HMC). This includes tasks such as satisfying adapter requirements and operating system requirements for partition mobility.

Software applications that recognize partition mobility:

Software applications might be designed to recognize and adapt to changes in the system hardware after being moved from one system to another.

Most software applications running in Linux logical partitions will not require any changes to work correctly during active partition mobility. Some applications might have dependencies on characteristics that change between the source and destination servers and other applications might need to adjust to support the migration.

PowerHA[®] (or High Availability Cluster Multi-Processing) clustering software is aware of partition mobility. You can move a mobile partition that is running the PowerHA clustering software to another server without restarting the PowerHA clustering software.

Examples of applications that would benefit if they were aware of partition mobility:

- Software applications that use processor and memory affinity characteristics to tune their behavior because affinity characteristics might change as a result of migration. The application's functions remain the same, but performance variations may be observed.
- Applications that use processor binding will maintain their binding to the same logical processors across migrations, but in reality the physical processors will change. Binding is usually done to maintain hot caches, but the physical processor move operation will require a cache hierarchy on the destination system. This usually occurs very quickly and should not be visible to the users.
- Applications that are tuned for given cache architectures, such as hierarchy, size, line-size, and associativity. These applications are usually limited to high-performance computing applications, but the just-in-time (JIT) compiler of the Java Virtual Machine is also optimized for the cache-line size of the processor on which it was opened.
- Performance analysis, capacity planning, and accounting tools and their agents are usually migration-aware because the processor performance counters might change between the source and destination servers, as might the processor type and frequency. Additionally, tools that calculate an aggregate system load based on the sum of the loads in all hosted logical partitions must be aware that a logical partition has left the system or that a new logical partition arrived.
- Workload managers

Network configuration in a partition mobility environment:

In partition mobility that is managed by the Hardware Management Console (HMC), the network between the source and destination servers is used to pass the mobile partition state information and other configuration data from the source environment to the destination environment. The mobile partition uses the virtual LAN for network access.

The virtual LAN must be bridged to a physical network using a Shared Ethernet Adapter in the Virtual I/O Server (VIOS) logical partition. The LAN must be configured so that the mobile partition can continue to communicate with other necessary clients and servers after a migration is completed.

Active partition mobility has no specific requirements on the memory size of the mobile partition or the type of network that connects the mover service partitions. The memory transfer does not interrupt the activity of the mobile partition. The memory transfer might take time when a large memory configuration

is busy on a slow network. Therefore, you might want to use a high-bandwidth connection, such as Gigabit Ethernet or faster, between the mover service partitions. The network bandwidth between the mover service partitions must be 1 GB or greater.

With VIOS 2.1.2.0, or later, you can enable secure IP tunnels between the mover service partition on the source server and the mover service partition on the destination server. For example, you might want to enable secure IP tunnels when the source and destination servers are not on a trusted network. Secure IP tunnels encrypt the partition state information that the mover service partitions exchange during active partition mobility. Mover service partitions with secure IP tunnels might require slightly more processing resources.

The Shared Ethernet Adapter bridges internal virtual LANs on the system with the external network such as the checkpoint firewall. With VIOS 2.2.1.4, or later, you can use the Trusted Firewall feature that is supported on the PowerSC™ Editions. With the Trusted Firewall capability, you can perform intervirtual LAN routing functions by using the Security Virtual Machine (SVM) kernel extension. By using this function, mobile partitions that are present on different virtual LANs of the same server can communicate by using the Shared Ethernet Adapter. During partition mobility, the SVM kernel extension checks for notification of network resumption on a migrated logical partition.

The maximum distance between the source and destination servers is decided by the following factors:

- The network and storage configuration that is used by the servers
- The ability of the applications to continue to operate when their storage is separated from the server by such a distance

If both servers are on the same network and are connected to the same shared storage, then active partition mobility validation succeeds. The time it takes to move the mobile partition, and the application performance after a move across a long distance, is dependent on the following factors:

- The network distance between the source and destination servers
- Application sensitivity to increased storage latency

Related concepts:

“Source and destination Virtual I/O Server logical partitions in a partition mobility environment” on page 31

Partition mobility that is managed by a Hardware Management Console (HMC) requires at least one Virtual I/O Server (VIOS) logical partition on the source server and at least one VIOS logical partition on the destination server.

Related tasks:

“Preparing the network configuration for partition mobility” on page 71

You need to verify that the network configuration is configured correctly so that you can successfully move the mobile partition from the source server to the destination server by using the Hardware Management Console (HMC). This includes tasks such as creating a Shared Ethernet Adapter on the source and destination Virtual I/O Server (VIOS) logical partitions and creating at least one virtual Ethernet adapter on the mobile partition.

Related information:

 [Trusted Firewall concepts](#)

Storage configuration in a partition mobility environment:

Learn about the virtual SCSI and virtual Fibre Channel configuration required for partition mobility that is managed by the Hardware Management Console (HMC).

Related concepts:

“Source and destination Virtual I/O Server logical partitions in a partition mobility environment” on page 31

Partition mobility that is managed by a Hardware Management Console (HMC) requires at least one Virtual I/O Server (VIOS) logical partition on the source server and at least one VIOS logical partition on the destination server.

Related tasks:


“Preparing the virtual SCSI configuration for partition mobility” on page 74

You need to verify that the virtual SCSI configuration is configured correctly so that you can successfully move the mobile partition from the source server to the destination server by using the Hardware Management Console (HMC). This includes tasks such as verifying the `reserve_policy` of the physical volumes, and verifying that the virtual devices have the same unique identifier, physical identifier, or IEEE volume attribute.

“Preparing the virtual Fiber Channel configuration for partition mobility” on page 79

You need to verify that the virtual Fiber Channel configuration is configured correctly so that you can successfully move the mobile partition from the source server to the destination server by using the Hardware Management Console (HMC).


Related information:

 [Virtual fibre channel](#)

Basic storage configuration in a partition mobility environment:

The mobile partition moves from one server to another by the source server sending the logical partition state information to the destination server over a local area network (LAN). However, partition disk data cannot pass from one system to another system over a network. Thus, for partition mobility to succeed, the mobile partition must use storage resources that are managed by a storage area network (SAN). By using SAN storage, the mobile partition can access the same storage from both the source and destination servers.

The following figure shows an example of the storage configuration required for partition mobility.



The physical storage that the mobile partition uses, Physical storage 3, is connected to the SAN. At least one physical adapter that is assigned to the source Virtual I/O Server logical partition is connected to the SAN. Similarly, at least one physical adapter that is assigned to the destination Virtual I/O Server logical partition is also connected to the SAN.

If the mobile partition connects to Physical storage 3 through virtual Fibre Channel adapters, the physical adapters that are assigned to the source and destination Virtual I/O Server logical partitions must support N_Port ID Virtualization (NPIV).

The mobile partition can use virtual I/O resources that are provided by one or more Virtual I/O Server logical partitions on the source server. To ensure successful mobility, configure the same number of Virtual I/O Server logical partitions on the destination server as are configured on the source server.

The physical adapter on the source Virtual I/O Server logical partition connects to one or more virtual adapters on the source Virtual I/O Server logical partition. Similarly, the physical adapter on the destination Virtual I/O Server logical partition connects to one or more virtual adapters on the destination Virtual I/O Server logical partition. If the mobile partition connects to Physical storage 3 through virtual SCSI adapters, the virtual adapters on both the source and destination Virtual I/O Server logical partitions are assigned to access the logical unit numbers (LUNs) of Physical storage 3.

Each virtual adapter on the source Virtual I/O Server logical partition connects to at least one virtual adapter on a client logical partition. Similarly, each virtual adapter on the destination Virtual I/O Server logical partition connects to at least one virtual adapter on a client logical partition.

Each virtual Fibre Channel adapter that is created on the mobile partition (or any client logical partition) is assigned a pair of worldwide port names (WWPNs). Both WWPNs in the WWPN pair are assigned to access the LUNs of the physical storage that the mobile partition uses, or Physical storage 3. During normal operation, the mobile partition uses one WWPN to log on to the SAN and access Physical Storage 3. When you move the mobile partition to the destination server, there is a brief period where the mobile partition runs on both the source and destination servers. Because the mobile partition cannot log on to the SAN from both the source and destination servers at the same time using the same WWPN, the mobile partition uses the second WWPN to log on to the SAN from the destination server during the migration. The WWPNs of each virtual Fibre Channel adapter move with the mobile partition to the destination server.

When you move the mobile partition to the destination server, the HMC (that manages the destination server) performs the following tasks on the destination server:

- Creates virtual adapters on the destination Virtual I/O Server logical partition
- Connects the virtual adapters on the destination Virtual I/O Server logical partition to the virtual adapters on the mobile partition

Redundancy configurations in a partition mobility environment:

In some situations, you can move a logical partition to a destination system that provides less redundancy than the source system.

The mobile partition can access its physical storage through redundant paths on the source system. The redundant paths can include redundant Virtual I/O Server (VIOS) logical partitions, VIOS logical partitions with redundant physical adapters, or both. In most cases, successful partition mobility requires that you maintain the same level of redundancy on the destination system as on the source system. Maintaining redundancy requires that you configure the same number of VIOS logical partitions with the same number of physical adapters on the source and destination servers.

In some situations, however, you might need to move a logical partition to a destination system with less redundancy than the source system. In these situations, you receive an error message explaining that the redundant configuration on the source system cannot be maintained on the destination system. Before you move the mobile partition, you can respond to the error in one of the following ways:


- You can change the configuration on the destination system so that you maintain redundancy.
- You can override virtual storage errors when possible. In other words, you can accept the reduced level of redundancy and continue with partition mobility.

The following table explains the configurations in which you can move a logical partition to a destination system with less redundancy than the source system. Some of these situations result in one or more failed paths to the physical storage after the mobile partition moves to the destination system.

Table 18. Redundancy options for partition mobility

Redundancy change	Source system	Destination system
<p>Redundant paths to the physical storage are maintained. However, the paths go through separate VIOS partitions on the source system and go through the same VIOS partition on the destination system.</p>	<p>The source system has two VIOS partitions. One physical Fibre Channel adapter in each VIOS partition provides the mobile partition with redundant paths to its physical storage.</p>	<p>The destination system has two VIOS partitions. Two physical Fibre Channel adapters in the VIOS partition provide the mobile partition with redundant paths to its physical storage.</p>
<p>Redundant paths to the physical storage are not maintained, and redundant VIOS partitions are not maintained. The mobile partition accesses its physical storage through redundant paths on the source system and through one path on the destination system.</p>	<p>The source system has two VIOS partitions. One physical adapter in each VIOS partition provides the mobile partition with redundant paths to its physical storage. (The physical and virtual adapters can be SCSI or Fibre Channel adapters.)</p>	<p>The destination system has one VIOS partition. One physical adapter in the VIOS partition provides the mobile partition with one path to its physical storage. (The physical and virtual adapters can be SCSI or Fibre Channel adapters.)</p> <p>This situation results in one successful path and one failed path to the physical storage. In an attempt to maintain redundancy, partition mobility creates two sets of virtual adapters. It maps one set of virtual adapters to the physical adapter, but it cannot map the other set of virtual adapters. The unmapped connections become a failed path.</p> <p>The paths consist of the following mappings. The adapters are either all SCSI adapters or all Fibre Channel adapters:</p> <ul style="list-style-type: none"> • The path to the physical storage consists of the following mappings: <ul style="list-style-type: none"> – A virtual client adapter to a virtual server adapter. – The virtual server adapter to the physical adapter. – The physical adapter to the physical storage. • The failed path consists of a virtual client adapter that is mapped to a virtual server adapter.
<p>Redundant paths to the physical storage are not maintained. The mobile partition accesses its physical storage through redundant paths on the source system and through one path on the destination system.</p>	<p>The source system has one VIOS partition. Two physical Fibre Channel adapters in the VIOS partition provide the mobile partition with redundant paths to its physical storage.</p>	<p>The destination system has one VIOS partition. One physical Fibre Channel adapter in the VIOS partition provides the mobile partition with one path to its physical storage.</p> <p>This situation results in one successful path and one failed path to the physical storage. In an attempt to maintain redundancy, partition mobility creates two sets of virtual adapters. It maps one set of virtual adapters to the physical adapter, but it cannot map the other set of virtual adapters. The unmapped connections become a failed path.</p>

Related information:

 Redundancy configuration using virtual fibre channel adapters

Preparing for partition mobility

You need to verify that the source and destination systems are configured correctly so that you can successfully move the mobile partition from the source system to the destination system. This includes verifying the configuration of the source and destination servers, the Hardware Management Console (HMC), the Virtual I/O Server logical partitions, the mobile partition, the virtual storage configuration, and the virtual network configuration.

Related concepts:

“Partition mobility overview for HMC” on page 4

You can learn about the benefits of partition mobility, how the Hardware Management Console (HMC) performs active and inactive partition mobility, and about the configuration that is required to successfully move a logical partition from one system to another.

“Partition mobility environment” on page 29

You can learn about each component of the partition mobility environment and its contribution in enabling successful partition mobility. Components of the partition mobility environment include the source and destination servers, the Hardware Management Console (HMC), the source and destination Virtual I/O Server logical partitions, the mobile partition, the networking configuration, and the storage configuration.

HMC-managed systems: Preparing the source and destination servers for partition mobility

You need to verify that the source and destination servers are configured correctly so that you can successfully move the mobile partition from the source server to the destination server by using the Hardware Management Console (HMC). This includes tasks such as verifying the logical memory block size of the source and destination servers, and verifying the available memory and processor resources of the destination server.

To prepare the source and destination servers for active or inactive partition mobility, complete the following tasks.

Table 19. Preparation tasks for the source and destination servers

Server planning tasks	Active mobility task	Inactive mobility task	Information resources
1. Ensure that the PowerVM hardware feature for IBM PowerLinux servers is activated.	X	X	<ul style="list-style-type: none">Entering the activation code for IBM PowerVM for IBM PowerLinux using the HMC version 7
2. If you do not have the PowerVM hardware feature for IBM PowerLinux servers, you can evaluate Live Partition Mobility at no cost by using Trial Live Partition Mobility. Ensure that you enter the activation code for Trial Live Partition Mobility.	X	X	<ul style="list-style-type: none">Entering the activation code for PowerVM for IBM PowerLinux using the HMC version 7

Table 19. Preparation tasks for the source and destination servers (continued)

Server planning tasks	Active mobility task	Inactive mobility task	Information resources
<p>3. Ensure that the source and destination servers are one of the following POWER7 models:</p> <ul style="list-style-type: none"> • 8246-L1C • 8246-L1D • 8246-L1S • 8246-L1T • 8246-L2C • 8246-L2D • 8246-L2S • 8246-L2T <p>Notes:</p> <ul style="list-style-type: none"> • The source and destination servers can also be POWER6 processor-based servers. See “Processor compatibility mode definitions” on page 12 for processor compatibility mode information. • Ensure that the destination server has the necessary software licenses and support maintenance contracts. To verify the entitlements that are active on your servers, see the Entitled Software Support website. 	X	X	
<p>4. Ensure that the firmware levels on the source and destination servers are compatible.</p>	X	X	“Firmware support matrix for partition mobility” on page 51
<p>5. Ensure that the source and destination servers are managed by an HMC in one of the following ways:</p> <ul style="list-style-type: none"> • The source and destination servers are managed by the same HMC (or redundant HMC pair). • The source server is managed by one HMC and the destination server is managed by a different HMC. 	X	X	
<p>6. Ensure that the logical memory block size is the same on the source and destination servers.</p>	X	X	Changing the logical memory block size
<p>7. Ensure that the destination server is not running on battery power. If the destination server is running on battery power, return the server to its regular power source before moving a logical partition.</p>	X	X	
<p>8. If the mobile partition uses shared memory, ensure that the shared memory pool is created on the destination server.</p>	X	X	Configuring the shared memory pool
<p>9. Ensure that the destination server has enough available memory to support the mobile partition.</p>	X		<ul style="list-style-type: none"> • If the mobile partition uses dedicated memory, see “Determining the available physical memory on the destination server” on page 53. • If the mobile partition uses shared memory, see “Determining the available I/O entitled memory on the destination server” on page 53.

Table 19. Preparation tasks for the source and destination servers (continued)

Server planning tasks	Active mobility task	Inactive mobility task	Information resources
10. Ensure that the destination server has enough available processors to support the mobile partition.	X		“Determining available processors on the destination server” on page 56
11. Verify that the source and destination mover service partitions can communicate with each other.	X		
12. Optional: Define the partition profile policy for inactive partition mobility.		X	“Defining the partition profile policy for inactive partition mobility” on page 55
<p>13. If the mobile partition on the source server is suspend-capable, verify that the destination server also supports suspend-capable partitions. You must also verify that there is a minimum of one reserved storage device with a size that can be determined by running the <code>lsrsdevsize</code> command from the HMC command line.</p> <p>Before you can use a disk from the reserved storage pool, you might need to clear the first 4096 bytes of the disk when you see a message that indicates that the disk cannot be used. The disk might have stale data indicating that it is already in use by another partition on the managed system, or might actively be used by another managed system. You must verify with your system administrator whether the disk is currently being used. You must initialize the first 4096 bytes of the disk only after you have ensured that the disk is no longer in use and that you have corrected all of the configuration issues associated with using the disk.</p>	X	X	<ul style="list-style-type: none"> To verify that the destination server supports suspend-capable partitions, see “Verifying that the destination server supports suspend-capable partitions” on page 55. To verify that there is a minimum of one reserved storage device with a size of at least 110% of the maximum partition memory, see “Determining the reserved storage device size in the destination server” on page 55.
If the mobile partition on the source server is a shared processor partition, and configured with processing units to virtual processor ratio of less than 0.1 and greater than or equal to 0.05, verify whether the destination server supports the processor minimum entitlement of 0.05 processor per virtual processor. The source and destination servers must be POWER7 processor-based servers.	X	X	You can verify that the destination server supports the same configuration as the source server by checking the processor level hardware capabilities of the destination server. To check the processor level hardware capabilities, see “Verifying the processor-level hardware capabilities of the destination server” on page 55.
If the mobile partition has Single Root IO Virtualization (SR-IOV) logical ports, then that partition cannot be migrated to the destination server.	X	X	

Table 19. Preparation tasks for the source and destination servers (continued)

Server planning tasks	Active mobility task	Inactive mobility task	Information resources
If the mobile partition is using a virtual Ethernet adapter which is using a virtual switch that is in the VEPA mode, or the mobile partition is using a virtual Ethernet adapter with a VSI profile, then verify that the destination server also supports virtual server network (VSN).	X	X	<ul style="list-style-type: none"> To verify that the destination server is capable of VSN, see “Verifying that the destination server supports the virtual server network” on page 127. To determine the virtual Ethernet switch name on the destination server, see “Determining the virtual Ethernet switch name and mode in the destination server” on page 128.

Related concepts:

“Source and destination servers in a partition mobility environment” on page 29

Two servers are involved in partition mobility that is managed by a Hardware Management Console (HMC). The *source server* is the server from which you want to move the logical partition, and the *destination server* is the server to which you want to move the logical partition.

Firmware support matrix for partition mobility:

Ensure that the firmware levels on the source and destination servers are compatible before upgrading.

In the following table, you can see that the first column represent the firmware level you are migrating from, and the values in the top row represent the firmware level you are migrating to. For each combination, *blocked* entries are blocked by code from migrating. *Not supported* entries are not blocked from migrating, but are not supported by IBM. *Mobile* entries are eligible for migration.

Table 20. Firmware level

Migrating from firmware level	350_xxx, or later	710_xxx	720_xxx	730_xxx	740_xxx	760_xxx	763_xxx	770_xxx	773_xxx	780_xxx
340_039, or later	Mobile	Mobile	Mobile	Blocked	Blocked	Blocked	Blocked	Blocked	Blocked	Blocked
350_xxx, or later	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile
710_xxx	Mobile	Mobile	Mobile	Mobile	Mobile	Blocked	Blocked	Blocked	Blocked	Blocked
720_xxx	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Blocked	Blocked	Blocked
730_xxx	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile
740_xxx	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile
760_xxx	Mobile	Blocked	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile
763_xxx	Mobile	Blocked	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile
770_xxx	Mobile	Blocked	Blocked	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile
773_xxx	Mobile	Blocked	Blocked	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile
780_xxx	Mobile	Blocked	Blocked	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile

The following table shows the number of concurrent migrations that are supported per system. The corresponding minimum levels of firmware, Hardware Management Console (HMC), and Virtual I/O Server (VIOS) that are required are also shown.

Table 21. Concurrent migrations

Concurrent migrations per system	Firmware level	HMC level	VMControl	VIOS level	Maximum concurrent migrations per VIOS
4	All	All	All	All	4
8	All	Version 7 Release 7.4.0, Service Pack 1, with mandatory fix MH01302, or later	VMControl Version 1.1.2, or later	Version 2.2.0.11, Fix Pack 24, Service Pack 1, or later	4
16	Level 7.6, or later	Version 7 Release 7.6.0, or later	VMControl V2.4.2	Version 2.2.2.0	8

Restrictions:

- Firmware levels 7.2 and 7.3 are restricted to eight concurrent migrations.
- Certain applications such as clustered applications, high availability solutions, and similar applications have heartbeat timers, also referred to as Dead Man Switch (DMS) for node, network, and storage subsystems. If you are migrating these types of applications, you must not use the concurrent migration option as it increases the likelihood of a timeout. This is especially true on 1 GB network connections.
- You must not perform more than four concurrent migrations on a 1 GB network connection. With VIOS Version 2.2.2.0 or later, and a network connection that supports 10 GB or higher, you can run a maximum of eight concurrent migrations.
- From VIOS Version 2.2.2.0, or later, you must have more than one pair of VIOS partitions to support more than eight concurrent mobility operations.
- Systems that are managed by the Integrated Virtualization Manager (IVM) support up to 10 concurrent migrations.
- To support the migration of up to 16 active or suspended mobile partitions from the source server to a single or multiple destination servers, the source server must have at least two VIOS partitions that are configured as mover service partitions. Each mover service partition must support up to 8 concurrent partition migration operations. If all 16 partitions are to be migrated to the same destination server, then the destination server must have at least two mover service partitions configured, and each mover service partition must support up to 8 concurrent partition migration operations.
- When the configuration of the mover service partition on the source or destination server does not support 8 concurrent migrations, any migration operation that is started by using either the graphical user interface or the command line will fail when no concurrent mover service partition migration resource is available. You must then use the **migr1par** command from the command line with the **-p** parameter to specify a comma-separated list of logical partition names, or the **--id** parameter to specify a comma-separated list of logical partition IDs.
- You can migrate a group of logical partitions by using the **migr1par** command from the command line. To perform the migration operations, you must use **-p** parameter to specify a comma-separated list of logical partition names, or the **--id** parameter to specify a comma-separated list of logical partition IDs.
- You can run up to four concurrent Suspend/Resume operations.
- You cannot perform Live Partition Mobility that is both bidirectional and concurrent. For example:
 - When you are moving a mobile partition from the source server to the destination server, you cannot move another mobile partition from the destination server to the source server.

- When you are moving a mobile partition from the source server to the destination server, you cannot move another mobile partition from the destination server to some other server.

Determining the available physical memory on the destination server:

You can determine whether the destination server has enough physical memory available to support the mobile partition and then make more physical memory available, if necessary, by using the Hardware Management Console (HMC).

You must be a super administrator to perform this task.

To determine whether the destination server has enough physical memory available to support the mobile partition, complete the following steps from the HMC:

1. Identify the amount of physical memory that the mobile partition requires:
 - a. In the navigation pane, expand **Systems Management > Servers**.
 - b. Click the source server on which the mobile partition is located.
 - c. In the work pane, select the mobile partition.
 - d. From the Tasks menu, click **Properties**. The Partition Properties window is displayed.
 - e. Click the **Hardware** tab.
 - f. Click the **Memory** tab.
 - g. Record the dedicated minimum, assigned, and maximum memory settings.
 - h. Click **OK**.
2. Identify the amount of physical memory that is available on the destination server:
 - a. In the navigation pane, expand **Systems Management** and click **Servers**.
 - b. In the work pane, select the destination server to which you plan to move the mobile partition.
 - c. From the **Tasks** menu, click **Properties**.
 - d. Click the **Memory** tab.
 - e. Record the **Current memory available for partition usage**.
 - f. Click **OK**.
3. Compare the values from steps 1 and 2. If the destination server does not have enough physical memory available to support the mobile partition, you can add more available physical memory to the destination server by performing one or more of the following tasks:
 - Dynamically remove physical memory from logical partitions that use dedicated memory. For instructions, see Removing dedicated memory dynamically.
 - If the destination server is configured with a shared memory pool, dynamically remove physical memory from the shared memory pool. For instructions, see Changing the size of the shared memory pool.

Determining the available I/O entitled memory on the destination server:

You can determine whether the shared memory pool on the destination server has enough available memory to accommodate the I/O entitled memory required by the mobile partition. You can then allocate more physical memory to the shared memory pool, if necessary, by using the Hardware Management Console (HMC).

You must be a super administrator to perform this task.

To determine whether the shared memory pool on the destination server has enough available memory to accommodate the I/O entitled memory required by the mobile partition, complete the following steps from the HMC:

1. Identify the amount of I/O entitled memory that the mobile partition requires:

- a. In the navigation pane, expand **Systems Management > Servers**.
 - b. Click the source server on which the mobile partition is located.
 - c. In the work pane, select the mobile partition.
 - d. From the Tasks menu, click **Properties**. The Partition Properties window is displayed.
 - e. Click the **Hardware** tab.
 - f. Click the **Memory** tab.
 - g. Click **Memory Statistics**. The Memory Statistics panel is displayed.
 - h. Record the **Assigned I/O Entitled Memory**. This is the amount of I/O entitled memory that the mobile partition requires on the destination server.
2. Identify the amount of available physical memory in the shared memory pool on the destination server:
 - a. In the navigation pane, expand **Systems Management** and click **Servers**.
 - b. In the work pane, select the destination server to which you plan to move the mobile partition.
 - c. From the **Tasks** menu, click **Configuration > Virtual Resources > Shared Memory Pool Management**.
 - d. Record the **Available pool memory** and click **OK**.
3. Compare the amount of available memory (from step 2) with the amount of I/O entitled memory required by the mobile partition (from step 1).
 - If more memory is available than the amount of I/O entitled memory required by the mobile partition, the shared memory pool on the destination server has enough available memory to support the mobile partition on the destination server.
 - If the amount of I/O entitled memory required by the mobile partition is greater than the amount of available memory, perform one or more of the following tasks:
 - Add memory to the shared memory pool so that the shared memory pool has enough available memory to accommodate the I/O entitled memory required by the mobile partition. For instructions, see *Changing the size of the shared memory pool*.
 - Remove one or more shared memory partitions from the shared memory pool until the shared memory pool has enough available memory to accommodate the I/O entitled memory required by the mobile partition. You can remove a logical partition from the shared memory pool by changing the memory mode of the logical partition from shared to dedicated. For instructions, see *Changing the memory mode of a logical partition*.
 - Remove I/O adapters from the mobile partition so that it requires less memory for I/O operations. For instructions, see *Removing virtual adapters dynamically*.
 - If the amount of I/O entitled memory that is required by the mobile partition is equal to, or almost equal to, the amount of available memory, the shared memory pool is probably greatly overcommitted, which can affect performance. Consider adding more memory to the shared memory pool to reduce the degree to which the shared memory pool is overcommitted.

Attention: If you move an active logical partition whose I/O entitled memory mode is set to auto, the HMC does not automatically recalculate and reassign the I/O entitled memory for the mobile partition until you restart the mobile partition on the destination server. If you restart the mobile partition on the destination server and you plan to move the mobile partition back to the source server, you must verify that the shared memory pool on the source server has enough available memory to accommodate the new amount of I/O entitled memory required by the mobile partition.

Related information:

 Performance considerations for overcommitted shared memory partitions

Defining the partition profile policy for inactive partition mobility:

You can select the partition profile policy for inactive partition mobility in the Hardware Management Console (HMC). You can either select the partition state defined in the hypervisor, or select the configuration data defined in the last activated profile on the source server. By default, the partition state defined in the hypervisor is selected.

To define a policy for inactive partition mobility, complete the following tasks:

1. In the navigation pane, open **Systems Management** and select **Servers**.
2. In the work pane, select the source server.
3. From the **Tasks** menu, select **Properties**.
4. Click the **Migration** tab.
 - To use the partition state that is defined in the hypervisor for memory and processor-related settings, select **Partition Configuration** in the **Inactive profile migration policy** list. However, if you are unable to start the partition, the data defined in the last activated profile on the source server is used, even though you select the **Partition Configuration** option.
 - To use the data defined in the last activated profile on the source-managed system for memory and processor-related settings, select **Last Activated Profile** in the **Inactive profile migration policy** list.
5. Click **OK**.

Verifying that the destination server supports suspend-capable partitions:

To move a Linux mobile partition that is capable of suspension, verify that the destination server supports suspend-capable partitions by using the Hardware Management Console (HMC).

With HMC 7.7.2.0, or later, you can suspend a Linux logical partition with its operating system and applications, and store its virtual server state to persistent storage. At a later stage, you can resume the operation of the logical partition. To verify that the destination server supports suspend-capable partitions, complete the following tasks:

1. In the navigation pane, open **Systems Management** and select **Servers**.
2. Select the destination server in the work pane.
3. From the **Tasks** menu, select **Properties**.
4. Click the **Capabilities** tab.
 - If Partition Suspend Capable is **True**, the destination server supports suspend-capable partitions.
 - If Partition Suspend Capable is **False**, the destination server does not support suspend-capable partitions, and you cannot move the mobile partition to the server. To move the mobile partition, change the partition configuration so that it is not capable of suspension.
5. Click **OK**.

Determining the reserved storage device size in the destination server:

To ensure that you can perform the suspend operation on partitions that are capable of suspension, you must determine the size of the storage device on the destination server. The size is based on several configuration attributes. You can run the **lsrsdevsize** command from the HMC command line to determine the size of the storage device on the destination server.

Verifying the processor-level hardware capabilities of the destination server:

On POWER7 processor-based servers, to move a shared processor mobile partition that is configured with processing units to a virtual processor ratio of less than 0.1 and greater than or equal to 0.05, verify that the destination server supports the same configuration by checking the processor-level hardware capabilities of the destination server.

By reducing the minimum entitlement to 0.05 processing units per virtual processor for all logical partitions without physical I/O devices, you can create up to 20 partitions on a single physical processor.

To verify the processor level hardware capabilities of the destination server, run the following command from the Hardware Management Console (HMC) command-line interface:

```
lshwres -r proc -m vrm113-fsp --level sys
```

If the value of the `min_proc_units_per_virtual_proc` attribute is 0.05, the destination server has the same processor-level hardware capabilities as the source server.

Determining available processors on the destination server:

You can determine the available processors on the destination server and allocate more processors, if necessary, by using the Hardware Management Console (HMC).

You must be a super administrator to perform this task.

To determine the available processors on the destination server using the HMC, complete the following steps:

1. Determine how many processors the mobile partition requires:
 - a. In the navigation pane, open **Systems Management** and select **Servers**.
 - b. Select the managed server of your choice in the navigation pane.
 - c. In the work pane, select the logical partition of your choice.
 - d. Select **Properties** and select the **Hardware** tab and the **Processors** tab.
 - e. View the Processor section and record the minimum, maximum, and available processor settings.
 - f. Click **OK**.
2. Determine the processors available on the destination server:
 - a. In the navigation pane, open **Systems Management** and select **Servers**.
 - b. Select the managed server of your choice in the navigation pane.
 - c. Select **Properties** and the **Processors** tab.
 - d. Record the **Available processors**.
 - e. Click **OK**.
3. Compare the values from steps 1 and 2.
 - If the destination server has enough available processors to support the mobile partition, continue to “HMC-managed systems: Preparing the source and destination servers for partition mobility” on page 48.
 - If the destination server does not have enough available processors to support the mobile partition, use the HMC, to dynamically remove the processors from the logical partition or you can remove processors from logical partitions on the destination server.

Improving the performance of partition mobility:

You can improve the partition mobility performance by installing the latest available firmware, Hardware Management Console (HMC), and the Virtual I/O Server (VIOS) software on both the source and target mover service partitions.

The following table describes the VIOS processing unit resources that are suggested in addition to the resources already assigned to the VIOS to handle the existing virtual I/O resource requirements, when you are using a 10 GB network adapter for partition mobility.

Table 22. VIOS processing unit resources and concurrent migration requirements

	POWER7		POWER7+™	
	Dedicated processing units	Shared processor virtual processing units	Dedicated processing units	Shared processor virtual processing units
Single Migration	3	3	2	2
Up to 16 concurrent migrations	4	4	3	3

An increase in the processor entitlement might be required to increase the number of virtual processing units. If you are using a 1 GB network adapter, or if the bandwidth of the 10 GB network adapter links that are to be used for partition mobility reach peak levels (near 100% utilization), one more POWER7 or POWER7+ processing unit or virtual processing unit is recommended for improving the performance, regardless of the number of concurrent migrations.

You can use 1 GB of extra memory when you use a 10 GB network adapter for partition mobility. You do not require extra memory when you are using 1 GB network adapter.

The following are the suggested settings:

- Enable the *Large Send* and *Large Receive Offload* options on all network devices that are involved in partition mobility.
- Set `tcp_sendspace = 524288` and `tcp_recvspace = 524288`.
- Enable the Jumbo Frames option if the environment supports the option.

In some instances, few operating system levels of the client partition allows the partition state to be transferred more efficiently during the migration process. The minimum AIX® level that supports the transfer of a partition state is AIX 6.1 Technology Level 4.

The following AIX levels are recommended for the client partitions:

- AIX 6.1 Technology Level 8.
- AIX 7.1 Technology Level 2.

Server evacuation:

You can perform a server evacuation operation by using the Hardware Management Console (HMC) that is at Version 7 Release 7.8.0, or later. A server evacuation operation is used to move all migration capable logical partitions from one system to another. Any upgrade or maintenance operations can be performed after all the partitions are migrated and the source system is powered off.

You can migrate all the migration capable Linux partitions from the source server to the destination server by running the following command from the HMC command line:

```
migr1par -o m -m srcCec -t dstCec --all
```

Note: The following conditions apply for a partition that is considered as migration capable:

- The source server must not have any inbound or outbound migration operations that are in progress.
- The destination server must not have any outbound migration operations that are in progress.
- The HMC must be at Version 7 Release 7.8.0, or later.

To stop the migration of all the migration capable Linux partitions, run the following command from the HMC command line:

```
migr1par -o s -m srcCec --all
```

Preparing the HMC for partition mobility

You need to verify that the Hardware Management Console (HMC) that manage the source and destination servers are configured correctly so that you can move the mobile partition from the source server to the destination server.

To prepare the HMC or HMCs for active or inactive partition mobility, complete the following tasks.

Table 23. Preparation tasks for the HMC

HMC planning tasks	Active mobility task	Inactive mobility task	Information resources
<p>1. Ensure that the HMC that manages the source server and the HMC that manages the destination server meet the following version requirements:</p> <ul style="list-style-type: none"> • If the source server, the destination server, or both servers are POWER7 processor-based servers, ensure that the HMC or HMCs that manage the servers are at Version 7 Release 7.1, or later. • If the source server or the destination server is a POWER6 processor-based server, ensure that the HMC that manages that server is at Version 7 Release 3.5, or later. 	X	X	<ul style="list-style-type: none"> • Determining your HMC machine code version and release • Updating your HMC software
<p>2. If the source server is managed by one HMC and the destination server is managed by a different HMC, verify that the secure shell (SSH) authentication keys are set up correctly between the HMC that manages the source server and the HMC that manages the destination server.</p>	X	X	<p>“Verifying SSH authentication between the source and destination HMCs”</p>
<p>3. If the mobile partition on the source server is capable of suspension, ensure that the HMC that manages the destination server is at Version 7 Release 7.2, or later.</p>	X	X	<ul style="list-style-type: none"> • “Verifying that the destination server supports suspend-capable partitions” on page 55 • “Determining the reserved storage device size in the destination server” on page 55
<p>If the mobile partition on the source server is configured with less than 0.1 and greater than or equal to 0.05 processing units, ensure that the destination server supports the same configuration. The HMC must be at Version 7 Release 7.6.0, or later.</p>	X	X	<p>“Verifying the processor-level hardware capabilities of the destination server” on page 55</p>
<p>If the mobile partition on the source server uses the virtual server network (VSN), verify that the destination server also uses VSN. The HMC must be at Version 7 Release 7.7.0, or later.</p>	X	X	<p>“Verifying that the destination server supports the virtual server network” on page 127</p>

Related concepts:

“Hardware Management Console in a partition mobility environment” on page 30

Learn about the Hardware Management Console (HMC) and how you can use its Partition Migration wizard to move an active or inactive logical partition from one server to another server.

Verifying SSH authentication between the source and destination HMCs:

You can run the **mkauthkeys** command from the Hardware Management Console (HMC) that manages the source server to verify that the secure shell (SSH) authentication keys are correctly set up between the

HMC that manages the source server and the HMC that manages the destination server. SSH authentication allows the HMCs to send and receive partition mobility commands to and from each other.

To verify that the SSH authentication keys are set up correctly between the HMC that manages the source server and the HMC that manages the destination server, complete the following steps:

1. Run the following command from the HMC command line of the HMC that manages the source server:

```
mkauthkeys -u <remoteUserName> --ip <remoteHostName> --test
```

Where:

- *remoteUserName* is the name of the user on the HMC that manages the destination server. This parameter is optional. If you do not specify a user name for the HMC that manages the destination server, then the migration process uses the current user name as the *remoteUserName*.
- *remoteHostName* is the IP address or the host name of the HMC that manages the destination server.

If this command produces a return code of 0, then the SSH authentication keys are set up correctly between the HMC that manages the source server and the HMC that manages the destination server.

If this command produces an error code, then continue to the next step to set up the SSH authentication keys between the HMC that manages the source server and the HMC that manages the destination server.

2. Run the following command to set up the SSH authentication keys between the HMC that manages the source server and the HMC that manages the destination server:

```
mkauthkeys -u <remoteUserName> --ip <remoteHostName> -g
```

Where *remoteUserName* and *remoteHostName* represent the same values that they represented in the previous step.

The `-g` option automatically sets up the SSH authentication keys from the HMC that manages the source server to the HMC that manages the destination server, and it automatically sets up the SSH authentication keys from the HMC that manages the destination server to the HMC that manages the source server. If you do not include the `-g` option, the command automatically sets up the SSH authentication keys from the HMC that manages the source server to the HMC that manages the destination server, but the command does not automatically set up the SSH authentication keys from the HMC that manages the destination server to the HMC that manages the source server.

Preparing the source and destination Virtual I/O Server logical partitions for partition mobility

You must verify that the source and destination Virtual I/O Server (VIOS) logical partitions are configured correctly so that you can successfully move the mobile partition from the source server to the destination server by using the Hardware Management Console (HMC). This verification includes tasks such as verifying the version of the VIOS partitions and enabling the mover service partitions.

To prepare the source and destination VIOS partitions for active or inactive partition mobility, complete the following tasks.

Table 24. Preparation tasks for the source and destination VIOS partitions

VIOS planning tasks	Active mobility task	Inactive mobility task	Information resources
<p>1. Ensure that at least one VIOS partition is installed and activated on both the source and destination servers.</p> <p>If the mobile partition receives virtual storage resources from redundant VIOS partitions on the source server, install the same number of VIOS partitions on the destination server, if possible.</p> <p>Remember: In some situations, you can select the option to override virtual storage errors if possible and move a logical partition to a destination system with less redundancy.</p>	X	X	Installing the Virtual I/O Server and client logical partitions
<p>2. Ensure that the source and destination VIOS partitions are at the following versions:</p> <ul style="list-style-type: none"> • To move Linux logical partitions, ensure that the source and destination VIOS partitions are at Version 2.1.2.0, Service Pack 1, or later. <p>Notes:</p> <ul style="list-style-type: none"> • From VIOS Version 2.2.0.11, Fix Pack 24, Service Pack 1 to VIOS Version 2.2.1.0, Live Partition Mobility for a client partition that uses storage that is provisioned from a shared storage pool is not supported. • From VIOS Version 2.2.0.11, Fix Pack 24, Service Pack 1 to VIOS Version 2.2.2.2, the Suspend/Resume feature for a Linux logical partition, which uses storage that is exported from a VIOS partition that is backed up by a shared storage pool is not supported. • From VIOS Version 2.2.3.0, or later, you can import one or more physical volumes into a shared storage pool. <ul style="list-style-type: none"> – You must not move an active or inactive partition that can access the physical volume while the import operation is in progress. – You must deactivate the paging space device for a partition that supports the Suspend/Resume features by disabling the features in the partition profile. After the import operation completes, you can enable these features again before the partition profile is activated. 	X	X	<ul style="list-style-type: none"> • Virtual I/O Server and Integrated Virtualization Manager commands • Migrating the Virtual I/O Server • Updating the Virtual I/O Server
<p>3. Ensure that the mover service partition is enabled on one or more source and destination VIOS partitions.</p> <p>Note: From VIOS Version 2.2.0.11, Fix Pack 24, Service Pack 1 to VIOS Version 2.2.1.0, you cannot use a VIOS logical partition that uses a shared storage pool as a mover service partition.</p>	X		“Enabling the source and destination mover service partitions” on page 61

Table 24. Preparation tasks for the source and destination VIOS partitions (continued)

VIOS planning tasks	Active mobility task	Inactive mobility task	Information resources
<p>4. If the mobile partition uses shared memory, verify that at least one VIOS partition is assigned to the shared memory pool on the destination server (subsequently referred to as a <i>paging VIOS partition</i>) and that it is at release level 2.1.1, or later.</p> <p>If the mobile partition accesses its paging space device redundantly through two paging VIOS partitions and you want to maintain this redundancy on the destination server, verify that two paging VIOS partitions are assigned to the shared memory pool on the destination server.</p> <p>Notes:</p> <ul style="list-style-type: none"> • From VIOS Version 2.2.0.11, Fix Pack 24, Service Pack 1 to VIOS Version 2.2.1.0, you cannot use a VIOS logical partition that uses a shared storage pool as a paging space partition. • In VIOS Version 2.2.0.11, Fix Pack 24, Service Pack 1, or later, you cannot use logical units in shared storage pools as paging devices. 	X	X	<ul style="list-style-type: none"> • Configuring the shared memory pool • Adding a paging VIOS partition to the shared memory pool
<p>5. If the mobile partition uses shared memory, verify that the shared memory pool on the destination server contains a paging space device that satisfies the size requirements and redundancy configuration of the mobile partition.</p>	X	X	<p>“Verifying that the destination shared memory pool contains an available paging space device” on page 62</p>
<p>6. Optional: Synchronize the time-of-day clocks for the source and destination VIOS partitions.</p>	X		<p>Synchronizing the time-of-day clocks of the source and destination Virtual I/O Server logical partitions</p>

Related concepts:

“Source and destination Virtual I/O Server logical partitions in a partition mobility environment” on page 31

Partition mobility that is managed by a Hardware Management Console (HMC) requires at least one Virtual I/O Server (VIOS) logical partition on the source server and at least one VIOS logical partition on the destination server.

Enabling the source and destination mover service partitions:

You can enable the mover service partition attribute on a Virtual I/O Server logical partition by using the Hardware Management Console (HMC).

You must be a super administrator or operator to complete this task.

There must be at least one mover service partition on the source and destination servers for the mobile partition to participate in active partition mobility. If the mover service partition is disabled on either the source or the destination Virtual I/O Server (VIOS), the mobile partition can participate only in inactive partition mobility.

To enable the source and destination mover service partition using the HMC, complete the following steps:

1. In the navigation pane, open **Systems Management** and select **Servers**.

2. Select the managed server of your choice in the navigation pane.
3. In the work pane, select a VIOS logical partition and select **Properties**.
4. On the **General** tab, select **Mover Service Partition**, and click **OK**.
5. Repeat steps 3 and 4 for the destination server.

Verifying that the destination shared memory pool contains an available paging space device:

You can verify that the shared memory pool on the destination server contains a paging space device that satisfies the size requirements and redundancy configuration of the mobile partition by using the Hardware Management Console (HMC).

To verify that the shared memory pool on the destination server contains a paging space device that satisfies the size requirements and redundancy configuration of the mobile partition, complete the following steps from the HMC:

1. Identify the size requirements of the mobile partition. The paging space device for the Linux logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) must be at least the size of the maximum logical memory of the shared memory partition. To view the maximum logical memory of the mobile partition, complete the following steps:
 - a. In the navigation pane, expand **Systems Management > Servers**, and click the system on which the mobile partition is located.
 - b. In the work pane, select the mobile partition, click the **Tasks** button, and click **Properties**. The Partition Properties window is displayed.
 - c. Click the **Hardware** tab.
 - d. Click the **Memory** tab.
 - e. Record the maximum logical memory. This is the size requirement for the paging space device for the mobile partition.
2. Identify the redundancy configuration of the mobile partition. In the **Memory** tab of the Partition Properties of the mobile partition, record the number of Virtual I/O Server (VIOS) logical partitions (hereafter referred to as *paging VIOS partitions*) that are assigned to the mobile partition:
 - If the mobile partition is assigned a primary paging VIOS partition and no secondary paging VIOS partition is assigned, then the mobile partition does not use redundant paging VIOS partitions. In this case, the mobile partition uses a paging space device that can only be accessed by one paging VIOS partition in the shared memory pool.
 - If the mobile partition is assigned a primary paging VIOS partition and a secondary paging VIOS partition, then the mobile partition uses redundant paging VIOS partitions. In this case, the mobile partition uses a paging space device that can be accessed redundantly by both paging VIOS partitions in the shared memory pool.
3. View the paging space devices that are currently assigned to the shared memory pool on the destination server:
 - a. In the navigation pane, expand **Systems Management** and click **Servers**.
 - b. In the work pane, select the destination server.
 - c. From the **Tasks** menu, click **Configuration > Virtual Resources > Shared Memory Pool Management**. The Shared Memory Pool Management window is displayed.
 - d. Click the **Paging Devices** tab.
 - e. Take note of the available paging space devices, their size, and whether they are capable of redundancy.

Note: Paging space devices can only be assigned to one shared memory pool at a time. You cannot assign the same paging space device to shared memory pools on two different systems at the same time.

4. Determine whether the shared memory pool on the destination server has a suitable paging space device for the mobile partition.
 - a. If the mobile partition *does not use* redundant paging VIOS partitions, verify that there is an active paging space device that is not capable of redundancy and meets the size requirement of the mobile partition. If no such device exists, you have the following options:
 - You can add a paging space device to the shared memory pool on the destination server. For instructions, see Adding and removing paging space devices to and from the shared memory pool.
 - If the shared memory pool contains an available paging space device that satisfies the size requirements of the mobile partition, but is capable of redundancy, you can move the mobile partition to the destination server. In this case, when you move the mobile partition to the destination server (active partition mobility) or when you activate the mobile partition on the destination server (inactive partition mobility), the HMC assigns the paging space device that is capable of redundancy to the mobile partition.
 - b. If the mobile partition *uses* redundant paging VIOS partitions, verify that an active paging space device is capable of redundancy and meets the size requirements of the mobile partition. If no such device exists, you have the following options:
 - You can add a paging space device to the shared memory pool on the destination server. For instructions, see Adding and removing paging space devices to and from the shared memory pool.
 - If the shared memory pool contains an available paging space device that satisfies the size requirements of the mobile partition, but is not capable of redundancy, you can move the mobile partition to the destination server. When you move the mobile partition to the destination server (active partition mobility) or when you activate the mobile partition on the destination server (inactive partition mobility), the HMC assigns the paging space device that is not capable of redundancy to the mobile partition. However, instead of using redundant paging VIOS partitions on the destination server, the mobile partition uses only the paging VIOS partition that has access to the paging space device that is not capable of redundancy.

Related information:

 Paging space devices on systems that are managed by an HMC

Preparing the mobile partition for partition mobility

You need to verify that the mobile partition is configured correctly so that you can successfully move it from the source server to the destination server by using the Hardware Management Console (HMC). This includes tasks such as satisfying adapter requirements and operating system requirements for partition mobility.

To prepare the mobile partition for active or inactive partition mobility, complete the following tasks.

Table 25. Preparation tasks for the mobile partition

Mobile partition planning tasks	Active mobility task	Inactive mobility task	Information resources
1. Ensure that the operating system running in the mobile partition is the Linux operating system. Restriction: The mobile partition cannot be a Virtual I/O Server (VIOS) logical partition.	X	X	

Table 25. Preparation tasks for the mobile partition (continued)

Mobile partition planning tasks	Active mobility task	Inactive mobility task	Information resources
<p>2. Ensure that the operating system is at one of the following levels:</p> <ul style="list-style-type: none"> • Red Hat Enterprise Linux version 5 Update 5, or later • SUSE Linux Enterprise Server 10 Service Pack 3, or later • SUSE Linux Enterprise Server 11 Service Pack 1, or later <p>Earlier versions of the Linux operating system can participate in inactive partition mobility if the operating systems support virtual devices and POWER6 or POWER7 processor-based servers.</p>	X		
<p>3. Ensure that the DynamicRM tool package is installed.</p>	X		Service and productivity tools for Linux POWER servers
<p>4. Ensure that Resource Monitoring and Control (RMC) connections are established with the Linux mobile partition, the source and destination VIOS logical partitions, and the source and destination mover service partitions.</p>	X		“Verifying RMC connections for the mobile partition” on page 65
<p>5. Verify that the processor compatibility mode of the mobile partition is supported on the destination server.</p>	X	X	“Verifying the processor compatibility mode of the mobile partition” on page 66
<p>6. Ensure that the mobile partition is not enabled for redundant error path reporting.</p>	X	X	“Disabling the mobile partition for redundant error-path reporting” on page 67
<p>7. Ensure that the mobile partition is only using a virtual serial adapter for virtual terminal connections.</p>	X	X	“Disabling virtual serial adapters for the mobile partition” on page 67
<p>8. Ensure that the mobile partition is not part of a partition workload group.</p>	X	X	“Removing the mobile partition from a partition workload group” on page 68
<p>9. Ensure that the mobile partition is not using barrier synchronization register (BSR) arrays.</p>	X		“Disabling BSR arrays for the mobile partition” on page 68
<p>10. Ensure that the mobile partition has not specified a value for Minimum, Allocated or Maximum huge pages.</p>	X		“Disabling huge pages for the mobile partition” on page 69
<p>11. Ensure that the mobile partition does not have physical I/O adapters and single root I/O virtualization (SR-IOV) logical ports.</p>	X		<ul style="list-style-type: none"> • Moving physical I/O devices and slots dynamically • Removing physical I/O devices and slots dynamically • Removing a single root I/O virtualization logical port from a logical partition dynamically
<p>12. Ensure that the mobile partition does not use Host Ethernet Adapters (or Integrated Virtual Ethernet).</p>	X		“Removing logical Host Ethernet Adapters from the mobile partition” on page 70

Table 25. Preparation tasks for the mobile partition (continued)

Mobile partition planning tasks	Active mobility task	Inactive mobility task	Information resources
13. Optional: Determine the name of the partition profile for the mobile partition on the destination server.	X	X	
14. Ensure that the applications running in the mobile partition are mobility-safe or mobility-aware.	X		“Software applications that recognize partition mobility” on page 42
15. If you changed any partition profile attributes, shut down and activate the new profile for the new values to take effect.	X	X	Shutting down and restarting logical partitions

Related concepts:

“Mobile partition managed by an HMC in a partition mobility environment” on page 41

A *mobile partition* is a logical partition that you want to move from the source server to the destination server. You can move a running mobile partition, or active mobile partition, or you can move a powered off mobile partition, or inactive mobile partition, from the source server to the destination server.

Verifying RMC connections for the mobile partition:

You can verify the Resource Monitoring and Control (RMC) connection between the mobile partition and the Hardware Management Console (HMC). This RMC connection is required to perform active partition mobility.

You must be a super administrator to complete this task.

With RMC, you can configure response actions or scripts that manage general system conditions with little or no involvement from the system administrator. On the HMC, RMC is used as the main communication channel between Linux logical partitions and the HMC.

To verify an RMC connection for the mobile partition, complete the following steps:

1. Using the HMC command line, enter `lspartition -dlpar`.
 - If the results for your logical partition are <Active 1>, then the RMC connection is established. Skip the rest of this procedure and return to “Preparing the mobile partition for partition mobility” on page 63.
 - If the results for your logical partition are <Active 0> or your logical partition is not displayed in the command results, continue to the next step.
2. Verify that the RMC firewall port on the HMC is disabled.
 - If the RMC firewall port is disabled, skip to step 3.
 - If the RMC firewall port is enabled, change your HMC firewall setting. Repeat step 1.
3. Use telnet to access the logical partition. If you cannot use telnet, open a virtual terminal on the HMC to set up the network on the logical partition.
4. If the logical partition network has been set up correctly and there is still no RMC connection, verify that the RSCT fileset is installed.
 - If the RSCT fileset is installed, use telnet to the HMC from the logical partition to verify if the network is working correctly and that the firewall has been disabled. After verifying these tasks, repeat step 1. If you continue to have problems establishing an RMC connection for your mobile partition, contact your next level of support.

Important: It takes approximately five minutes for RMC connection to establish the connection after the network setup has been changed or after activating the logical partition.

Verifying the processor compatibility mode of the mobile partition:

You can use the Hardware Management Console (HMC) to determine whether the processor compatibility mode of the mobile partition is supported on the destination server, and update the mode, if necessary, so that you can successfully move the mobile partition to the destination server.

To verify that the processor compatibility mode of the mobile partition is supported on the destination server by using the HMC, complete the following steps:

1. Identify the processor compatibility modes that are supported by the destination server by entering the following command on the command line of the HMC that manages the destination server:

```
lssyscfg -r sys -F lpar_proc_compat_modes
```

Record these values so that you can refer to them later.

2. Identify the preferred processor compatibility mode of the mobile partition:
 - a. In the navigation pane of the HMC that manages the source server, open **Systems Management > Servers** and select the source server.
 - b. In the work pane, select the mobile partition.
 - c. From the Tasks menu, select **Configuration > Manage Profiles**. The Managed Profiles window is displayed.
 - d. Select the active partition profile of the mobile partition or select the partition profile from which the mobile partition was last activated.
 - e. From the Actions menu, click **Edit**. The Logical Partition Profile Properties window is displayed.
 - f. Click the **Processors** tab to view the preferred processor compatibility mode. Record this value so that you can refer to it later.
3. Identify the current processor compatibility mode of the mobile partition. If you plan to perform an inactive migration, skip this step and go to step 4.
 - a. In the navigation pane of the HMC that manages the source server, expand **Systems Management > Servers** and select the source server.
 - b. In the work pane, select the mobile partition and click **Properties**.
 - c. Select the **Hardware** tab and view the Processor Compatibility Mode. This is the current processor compatibility mode of the mobile partition. Record this value so that you can refer to it later.
4. Verify that the preferred and current processor compatibility modes that you identified in steps 2 and 3 are in the list of supported processor compatibility modes that you identified in step 1 for the destination server. For active migrations, both the preferred and current processor compatibility modes of the mobile partition must be supported by the destination server. For inactive migrations, only the preferred processor compatibility mode must be supported by the destination server.

Attention: If the current processor compatibility mode of the mobile partition is the POWER5 mode, be aware that the POWER5 mode does not appear in the list of modes supported by the destination server. However, the destination server supports the POWER5 mode even though it does not appear in the list of supported modes.
5. If the preferred processor compatibility mode of the mobile partition is not supported by the destination server, use step 2 to change the preferred mode to a mode that is supported by the destination server. For example, the preferred mode of the mobile partition is the POWER7 mode and you plan to move the mobile partition to a POWER6 processor-based server. The POWER6 processor-based server does not support the POWER7 mode, but it does support the POWER6 mode. Therefore, you change the preferred mode to the POWER6 mode.
6. If the current processor compatibility mode of the mobile partition is not supported by the destination server, try the following solutions:
 - If the mobile partition is active, it is possible that the hypervisor has not had the opportunity to update the current mode of the mobile partition. Restart the mobile partition so that the hypervisor can evaluate the configuration and update the current mode of the mobile partition.

- If the current mode of the mobile partition still does not match the list of supported modes that you identified for the destination server, use step 2 on page 66 to change the preferred mode of the mobile partition to a mode that is supported by the destination server.

Then, restart the mobile partition so that the hypervisor can evaluate the configuration and update the current mode of the mobile partition.

For example, assume that the mobile partition runs on a POWER7 processor-based server and its current mode is the POWER7 mode. You want to move the mobile partition to a POWER6 processor-based server, which does not support the POWER7 mode. You change the preferred mode of the mobile partition to the POWER6 mode and restart the mobile partition. The hypervisor evaluates the configuration and sets the current mode to the POWER6 mode, which is supported on the destination server.

Related concepts:

“Processor compatibility modes” on page 12

Processor compatibility modes enable you to move logical partitions between servers that have different processor types without upgrading the operating environments installed in the logical partitions.

“Processor compatibility modes” on page 99

Processor compatibility modes enable you to move logical partitions between servers that have different processor types without upgrading the operating environments installed in the logical partitions.

Disabling the mobile partition for redundant error-path reporting:

You can disable the mobile partition for redundant error-path reporting by using the Hardware Management Console (HMC) so that you can move the mobile partition from the source server to the destination server.

You must be a super administrator to complete this task.

If you enable redundant error-path reporting, the logical partition reports common server hardware errors and partition hardware errors to the HMC. If you disable redundant error-path reporting, the logical partition reports only partition hardware errors to the HMC. If you want to move a logical partition, disable the redundant error-path reporting.

To disable the mobile partition for redundant error-path reporting using the HMC, complete the following steps:

1. In the navigation pane, open **Systems Management** and select **Servers**.
2. Select the managed server of your choice in the navigation pane.
3. In the work pane, select the logical partition of your choice.
4. Select **Configuration > Manage Profiles**.
5. Select the profile of your choice and select **Actions > Edit**.
6. Click the **Settings** tab.
7. Deselect **Enable redundant error path reporting** and click **OK**. For this change to take effect, activate this logical partition with this profile.

Disabling virtual serial adapters for the mobile partition:

You can disable unreserved virtual serial adapters for the mobile partition by using the Hardware Management Console (HMC) so that you can move the mobile partition from the source server to the destination server.

You must be a super administrator to complete this task.

Virtual serial adapters are often used for virtual terminal connections to the operating system. The first two virtual serial adapters (slots 0 and 1) are reserved for the HMC. For a logical partition to participate in partition mobility, it cannot have any virtual serial adapters, except for the two that are reserved for the HMC.

To disable unreserved virtual serial adapters using the HMC, complete the following steps:

1. In the navigation pane, open **Systems Management** and select **Servers**.
2. Select the managed server of your choice in the navigation pane.
3. In the work pane, select the logical partition of your choice.
4. Select **Configuration > Manage Profiles**.
5. Select the profile of your choice and select **Actions > Edit**.
6. Select the **Virtual Adapter** tab.
7. If there are more than two virtual serial adapters listed, then ensure that the additional adapters beyond 0 and 1 are not selected as **Required**.
 - If you have additional virtual serial adapters listed as **Required**, select the adapter that you would like to remove. Then select **Actions > Delete** to remove the adapter from the partition profile.
 - You can select **Dynamic Logical Partitioning > Virtual Adapters**. The Virtual Adapters panel is displayed. Select the adapter that you would like to remove and select **Actions > Delete** to remove the adapter from the partition profile.
8. Click **OK**.

Removing the mobile partition from a partition workload group:

You can remove the mobile partition from a partition workload group by using the Hardware Management Console (HMC) so that you can move the mobile partition from the source server to the destination server.

You must be a super administrator to complete this task.

A partition workload group identifies a set of logical partitions that are located on the same physical system. The partition profile specifies the name of the partition workload group that it belongs to, if applicable. A partition workload group is defined when you use the HMC to configure a logical partition. For a logical partition to participate in partition mobility, it cannot be assigned to a partition workload group.

To remove the mobile partition from a partition workload group using the HMC, complete the following steps:

1. In the navigation pane, open **Systems Management** and select **Servers**.
2. Select the managed server of your choice in the navigation pane.
3. In the work pane, select the logical partition of your choice.
4. Select **Configuration > Manage Profiles**.
5. Select the profile of your choice and select **Actions > Edit**.
6. Click the **Settings** tab.
7. In the Workload Management area, select **(None)** and click **OK**.
8. Repeat steps 1 through 7 for all partition profiles associated with the mobile partition. For this change to take effect, you will need to activate this logical partition with this profile.

Disabling BSR arrays for the mobile partition:

You can disable barrier synchronization register (BSR) arrays for the mobile partition by using the Hardware Management Console (HMC) so that you can perform active partition mobility.

You must be a super administrator to perform this task.

BSR is a memory register that is located on certain POWER processor-based systems.

For a logical partition to participate in active partition mobility, it cannot use BSR arrays. If the mobile partition uses BSR, the logical partition can participate in inactive partition mobility.

To disable BSR for the mobile partition using the HMC, complete the following steps:

1. In the navigation pane, select **Systems Management** and select **Servers**.
2. In the navigation pane, select the managed server of your choice and select **Properties**.
3. Click the **Capabilities** tab.
 - If barrier synchronization register (BSR) Capable is **True**, click **OK** and continue with the next step.
 - If barrier synchronization register (BSR) Capable is **False**, the server does not support BSR. Skip the rest of this procedure and continue to “Preparing the mobile partition for partition mobility” on page 63.
4. In the navigation pane, open **Systems Management** and select **Servers**.
5. Select the managed server of your choice in the navigation pane.
6. In the work pane, select the logical partition of your choice, click the **Tasks** button, and select **Properties**.
7. Click the **Hardware** tab.
8. Click the **Memory** tab.
 - If the number of BSR arrays equals zero, the mobile partition can participate in active or inactive partition mobility. Skip the rest of this procedure and continue to “Preparing the mobile partition for partition mobility” on page 63.
 - If the number of BSR arrays is not equal to zero, then take one of the following actions:
 - Perform an inactive movement instead of an active movement.
 - Click **OK** and continue to the next step to prepare the mobile partition for an active movement.
9. Select the mobile partition, and then select **Configuration > Manage Profiles**.
10. Select the partition profile with which you will reactivate the mobile partition, and select **Action > Edit**.
11. Click the **Memory** tab.
 - If the number of BSR arrays equals 0, the mobile partition can participate in active or inactive partition mobility. Skip the rest of this procedure and continue to “Preparing the mobile partition for partition mobility” on page 63.
 - If the number of BSR arrays is not equal to 0, then take the following action to change BSR to 0 if you want to do an active migration:
 - Enter 0 in the field for the BSR arrays.
 - Click **OK** and continue to the next step to prepare the mobile partition for an active movement.
12. Activate this logical partition with this profile in order for this change to take effect.

Disabling huge pages for the mobile partition:

You can disable huge pages for the mobile partition by using the Hardware Management Console (HMC) so that you can perform active partition mobility.

You must be a super administrator to perform this task.

Huge pages can improve performance in specific environments that require a high degree of parallelism, such as in DB2 partitioned database environments. You can specify the minimum, desired, and maximum number of huge pages to assign to a logical partition when you create the logical partition or partition profile.

For a logical partition to participate in active partition mobility, it cannot use huge pages. If the mobile partition uses huge pages, it can participate in inactive partition mobility.

To disable huge pages for the mobile partition using the HMC, complete the following steps:

1. In the navigation pane, open **Systems Management** and select **Servers**.
2. In the work pane, select the managed server of your choice, click the **Tasks** button, and select **Properties**.
3. Click the **Capabilities** tab.
 - If Huge Page Capable is **True**, then click **OK** and continue with the next step.
 - If Huge Page Capable is **False**, then the source server does not support huge pages. The mobile partition can participate in active or inactive partition mobility. Skip the rest of this procedure and continue to “Preparing the mobile partition for partition mobility” on page 63.
4. In the navigation pane, open **Systems Management** and select **Servers**.
5. Select the managed server of your choice in the navigation pane.
6. In the work pane, select the logical partition of your choice.
7. Select **Properties** and the **Hardware** tab and then click the **Memory** tab.
 - If the **Minimum, Allocated** and **Maximum** huge page memory equals 0, then skip the rest of this procedure and continue to “Preparing the mobile partition for partition mobility” on page 63.
 - If the current huge page memory is not equal to 0, then take one of the following actions:
 - Perform an inactive movement instead of an active movement.
 - Click **OK** and continue with the next step to prepare the mobile partition for an active movement.
8. In the navigation pane, open **Systems Management** and select **Servers**.
9. Select the managed server of your choice in the navigation pane.
10. In the work pane, select the logical partition of your choice.
11. Select **Configuration > Manage Profiles**.
12. Select the profile of your choice and select **Actions > Edit**.
13. Click the **Memory** tab.
14. In the **Huge Page Memory** area, enter **0** in the **Minimum, Allocated** and **Maximum** fields and click **OK**.
15. Activate this logical partition with this profile in order for this change to take effect.

Removing logical Host Ethernet Adapters from the mobile partition:

You can remove a logical Host Ethernet Adapter (LHEA) from a mobile partition by using the Hardware Management Console (HMC) so that you can perform active partition mobility.

You must be a super administrator to perform this task.

For a logical partition to participate in active partition mobility, it cannot be assigned any LHEAs. If the mobile partition is assigned one or more LHEAs, it can participate in inactive partition mobility.

To remove an LHEA from the mobile partition using the HMC, complete the following steps:

1. In the navigation pane, open **Systems Management** and select **Servers**.
2. Select the managed server of your choice in the navigation pane.

3. Select the mobile partition and select **Configuration > Manage Profiles**.
4. Select the partition profile of your choice and select **Actions > Edit**.
5. Select the **Logical Host Ethernet Adapters (LHEA)** tab.
6. Select the physical port locations that have a logical port ID assigned to it and click **Reset**.
7. Click **OK**.

Preparing the network configuration for partition mobility

You need to verify that the network configuration is configured correctly so that you can successfully move the mobile partition from the source server to the destination server by using the Hardware Management Console (HMC). This includes tasks such as creating a Shared Ethernet Adapter on the source and destination Virtual I/O Server (VIOS) logical partitions and creating at least one virtual Ethernet adapter on the mobile partition.

To prepare the network configuration for active or inactive partition mobility, complete the following tasks.

Note: Partition mobility fails if you have enabled one of the following security settings on the VIOS logical partitions:

- If you have set network security to the high mode by using the **viosecure** command on the VIOS command-line interface
- If you have enabled a profile that impacts network connectivity by using the **viosecure** command on the VIOS command-line interface

You can enable secure IP tunnels between the mover service partitions on the source and destination servers to perform partition mobility with these security settings. For more information, see “Configuring secure IP tunnels between the mover service partitions on the source and destination servers” on page 72.

Table 26. Planning tasks for the network

Network planning tasks	Active mobility task	Inactive mobility task	Information resources
1. Create a Shared Ethernet Adapter on the source and destination Virtual I/O Server logical partition using the HMC.	X	X	Creating a Shared Ethernet Adapter for a VIOS logical partition using the HMC
2. Configure virtual Ethernet adapters on the source and destination Virtual I/O Server logical partitions.	X	X	Configuring a virtual Ethernet adapter using the HMC
3. Create at least one virtual Ethernet adapter on the mobile partition. Note: During a partition migration or suspend operation, if the source partition has at least one virtual Ethernet adapter that is disabled, the migration or suspend operation fails.	X		Configuring a virtual Ethernet adapter using the HMC
4. Activate the mobile partition to establish communication between the virtual Ethernet adapter and Virtual I/O Server virtual Ethernet adapter.	X		Activating a logical partition
5. Verify that the operating system of the mobile partition recognizes the new Ethernet adapter.	X		
6. Set up the LAN so that the mobile partition can continue to communicate with other necessary clients and servers after the migration is completed.	X	X	
7. Optional: Configure and enable secure IP tunnels between the mover service partitions on the source and destination servers.	X		“Configuring secure IP tunnels between the mover service partitions on the source and destination servers” on page 72

Table 26. Planning tasks for the network (continued)

Network planning tasks	Active mobility task	Inactive mobility task	Information resources
8. For VIOS partitions that are designated as mover service partitions, ensure that the network bandwidth between them is 1 GB or greater.	X		

Note:

- Partition mobility fails when the Virtual Station Interface (VSI) configuration on the destination server fails. You can use the `--vsi` override flag with the `migr1par` command to continue with the migration.
- Certain applications (like clustered applications, high availability solutions and other such applications) have heartbeat timers, also referred to as Dead Man Switch (DMS) for node, network, and storage subsystems. During partition mobility operations, there is normally a short period when the heartbeat function is suspended. The following are ways to reduce the likelihood of a heartbeat timeout:
 - When the line speed is higher, occurrence of a heartbeat timeout reduces. It is recommended to have a 10-Gigabit Ethernet connection on both the source and the destination system that is dedicated to Live Partition Mobility.
 - If you are running applications that are based on AIX, upgrade to AIX 6.1 Technology Level 8, or later, or AIX 7.1 Technology Level 2, or later.
 - Ensure that you are using the latest HMC and server firmware for the system.
 - Disable the heartbeat timer or lengthen the timeout value before starting the partition mobility operation, and re-enable the timer after the completion of the partition mobility operation.

Related concepts:

“Network configuration in a partition mobility environment” on page 42

In partition mobility that is managed by the Hardware Management Console (HMC), the network between the source and destination servers is used to pass the mobile partition state information and other configuration data from the source environment to the destination environment. The mobile partition uses the virtual LAN for network access.

Related information:

 [viosecure command](#)

Configuring secure IP tunnels between the mover service partitions on the source and destination servers:

With Virtual I/O Server (VIOS) 2.1.2.0, or later, you can configure secure IP tunnels between the mover service partitions on the source and destination servers. However, when both the source and destination servers are using the Virtual I/O Server 2.2.2.0, or later, the tunnels are created automatically depending on the security profile applied on the source VIOS.

Consider enabling secure IP tunnels between the mover service partition on the source server and the mover service partition on the destination server. For example, you might want to enable secure IP tunnels when the source and destination servers are not on a trusted network. Secure IP tunnels encrypt the partition state data that the mover service partition on the source server sends to the mover service partition on the destination server during active partition mobility.

Before you start, complete the following tasks:

1. Verify that the mover service partitions on the source and destination servers are at version 2.1.2.0, or later, by using the `ioslevel` command.
2. Obtain the IP address of the mover service partition on the source server.
3. Obtain the IP address of the mover service partition on the destination server.

4. Obtain the preshared authentication key for the source and destination mover service partitions.

To configure and enable secure IP tunnels, complete the following steps:

1. List the available secure tunnel agents by using the **lssvc** command. For example:

```
$lssvc  
ipsec_tunnel
```

2. List all the attributes that are associated with the secure tunnel agent by using the **cfgsvc** command.
For example:

```
$cfgsvc ipsec_tunnel -ls  
local_ip  
remote_ip  
key
```

3. Configure a secure tunnel between the mover service partition on the source server and the mover service partition on the destination server by using the **cfgsvc** command:

```
cfgsvc ipsec_tunnel -attr local_ip=src_msp_ip remote_ip=dest_msp_ip key=key
```

where:

- *src_msp_ip* is the IP address of the mover service partition on the source server.
 - *dest_msp_ip* is the IP address of the mover service partition on the destination server.
 - *key* is the preshared authentication key for the mover service partitions on the source and destination servers. For example, abcderadf31231adsf.
4. Enable the secure tunnel by using the **startsvc** command. For example:
startsvc ipsec_tunnel

Note: When you apply the High, Payment Card Industry (PCI), or Department of Defence (DoD) security profiles, the secure tunnel is created and active partition mobility is performed over this secure channel. The secure channel that was created automatically gets destroyed when the partition mobility operation is complete.

Related concepts:

“Source and destination Virtual I/O Server logical partitions in a partition mobility environment” on page 31

Partition mobility that is managed by a Hardware Management Console (HMC) requires at least one Virtual I/O Server (VIOS) logical partition on the source server and at least one VIOS logical partition on the destination server.

“Integrated Virtualization Manager in a partition mobility environment” on page 117

Learn about the Integrated Virtualization Manager (IVM) and how you can use it to move an active or inactive logical partition from one server to another server.


“Network configuration in a partition mobility environment” on page 42

In partition mobility that is managed by the Hardware Management Console (HMC), the network between the source and destination servers is used to pass the mobile partition state information and other configuration data from the source environment to the destination environment. The mobile partition uses the virtual LAN for network access.

“Network configuration in a partition mobility environment” on page 119

In partition mobility that is managed by the Integrated Virtualization Manager (IVM), the network between the source and destination servers is used to pass the mobile partition state information and other configuration data from the source environment to the destination environment. The mobile partition uses the virtual LAN for network access.

Related information:

 [cfgsvc command](#)

 [startsvc command](#)

Preparing the virtual SCSI configuration for partition mobility

You need to verify that the virtual SCSI configuration is configured correctly so that you can successfully move the mobile partition from the source server to the destination server by using the Hardware Management Console (HMC). This includes tasks such as verifying the reserve_policy of the physical volumes, and verifying that the virtual devices have the same unique identifier, physical identifier, or IEEE volume attribute.

The destination server must provide the same virtual SCSI configuration as the source server. In this configuration, the mobile partition can access its physical storage on the storage area network (SAN) after it moves to the destination server.

The Peer-to-Peer Remote Copy (PPRC) feature is supported on the virtual target device. The hardware-based disaster recovery solutions Global Mirror and Metro Mirror are based on PPRC. These solutions provide real-time mirroring of disks within an Enterprise Storage Server® or between two distant Enterprise Storage Servers.

To prepare the virtual SCSI configuration for active or inactive partition mobility, complete the following tasks.

Table 27. Preparation tasks for the virtual SCSI configuration on systems that are managed by the HMC

Storage planning tasks	Active mobility task	Inactive mobility task	Information resources
1. Verify that the physical storage that is used by the mobile partition is assigned to at least one Virtual I/O Server (VIOS) partition on the source server and to at least one VIOS partition on the destination server.	X	X	IBM System Storage® SAN Volume Controller
2. Verify that the reserve attributes on the physical volumes are the same for the source and destination VIOS partitions.	X	X	"Setting the reserve policy attributes of a device" on page 75
3. Verify that the virtual devices have the same unique identifier, physical identifier, or an IEEE volume attribute.	X	X	Identifying exportable disks
4. Verify that the virtual SCSI adapters on the mobile partition can access the virtual SCSI adapters on the source VIOS partition.	X	X	"Verifying the virtual adapter connections between the mobile partition and the Virtual I/O Server logical partitions on the source server" on page 76
5. Optional: Specify a new name for one or more virtual target devices to use on the destination VIOS partition.	X	X	"Specifying a new name for a virtual target device to use on a destination VIOS partition" on page 78
6. Verify that the mobile partition has access to the physical storage on the SAN.	X	X	"Verifying that the mobile partition has access to its physical storage" on page 77
7. If you changed any partition profile attributes, restart the mobile partition for the new values to take effect.	X	X	Shutting down and restarting logical partitions

Related concepts:

"Storage configuration in a partition mobility environment" on page 43

Learn about the virtual SCSI and virtual Fibre Channel configuration required for partition mobility that is managed by the Hardware Management Console (HMC).

Setting the reserve policy attributes of a device:

In some configurations, you must consider the reservation policy of the device on the Virtual I/O Server (VIOS).

The following table explains the situations in which the reservation policy of the device on the VIOS is important for systems that are managed by the Hardware Management Console (HMC) and the Integrated Virtualization Manager (IVM).

Table 28. Situations where the reservation policy of a device is important

HMC-managed systems	IVM-managed systems
<ul style="list-style-type: none"> To use a Multipath I/O (MPIO) configuration at the client, none of the virtual Small Computer Serial Interface (SCSI) devices on the VIOS can reserve the virtual SCSI device. Set the <code>reserve_policy</code> attribute of the device to <code>no_reserve</code>. For virtual SCSI devices used with Live Partition Mobility or the Suspend/Resume feature, the reserve attribute on the physical storage that is used by the mobile partition can be set as follows: <ul style="list-style-type: none"> You can set the reserve policy attribute to <code>no_reserve</code>. You can set the reserve policy attribute to <code>pr_shared</code> when the following products are at the following versions: <ul style="list-style-type: none"> HMC Version 7 release 3.5.0, or later VIOS Version 2.1.2.0, or later The physical adapters support the SCSI-3 Persistent Reserves standard <p>The reserve attribute must be the same on the source and destination VIOS partitions for successful partition mobility.</p> <ul style="list-style-type: none"> For PowerVM Active Memory™ Sharing or Suspend/Resume features, the VIOS automatically sets the reserve attribute on the physical volume to <code>no_reserve</code>. The VIOS performs this action when you add a paging space device to the shared memory pool. 	<p>For virtual SCSI devices used with Live Partition Mobility, the reserve attribute on the physical storage that is used by the mobile partition can be set as follows:</p> <ul style="list-style-type: none"> You can set the reserve policy attribute to <code>no_reserve</code>. You can set the reserve policy attribute to <code>pr_shared</code> when the following products are at the following versions: <ul style="list-style-type: none"> IVM Version 2.1.2.0, or later The physical adapters support the SCSI-3 Persistent Reserves standard <p>The reserve attribute must be the same on the source and destination management partitions for successful partition mobility.</p>

- From a VIOS partition, list the disks (or paging space devices) to which the VIOS has access. Run the following command:

```
lsdev -type disk
```

- To determine the reserve policy of a disk, run the following command, where `hdiskX` is the name of the disk that you identified in step 1. For example, `hdisk5`.

```
lsdev -dev hdiskX -attr reserve_policy
```

The results might look like the following output:

```
..
reserve_policy no_reserve                Reserve Policy                True
```

Based on the information in Table 28, you might need to change the `reserve_policy` so that you can use the disk in any of the described configurations.

- To set the `reserve_policy`, run the `chdev` command. For example:

```
chdev -dev hdiskX -attr reserve_policy=reservation
```

where:

- *hdiskX* is the name of the disk for which you want to set the `reserve_policy` attribute to `no_reserve`.
- *reservation* is either `no_reserve` or `pr_shared`.

4. Repeat this procedure from the other VIOS partition.

Requirements:

- Although the `reserve_policy` attribute is an attribute of the device, each VIOS saves the value of the attribute. You must set the `reserve_policy` attribute from both VIOS partitions so that both VIOS partitions recognize the `reserve_policy` of the device.
- For partition mobility, the `reserve_policy` on the destination VIOS partition must be the same as the `reserve_policy` on the source VIOS partition. For example, if the `reserve_policy` on the source VIOS partition is `pr_shared`, the `reserve_policy` on the destination VIOS partition must also be `pr_shared`.
- With the `PR_exclusive` mode on SCSI-3 reserve, you cannot migrate from one system to another system.
- The `PR_key` value for the VSCSI disks on the source system and the target system must be different.

Verifying the virtual adapter connections between the mobile partition and the Virtual I/O Server logical partitions on the source server:

You can verify the virtual adapter connections between the mobile partition and the Virtual I/O Server logical partitions on the source server so that the Hardware Management Console (HMC) correctly configures the virtual adapters on the destination server when you move the mobile partition.

To verify the virtual adapter connections between the mobile partition and the source Virtual I/O Server logical partitions, complete the following steps from the HMC:

- Verify the virtual adapter configuration of the mobile partition:
 - In the navigation pane, expand **Systems Management > Servers**.
 - Click the managed system on which the mobile partition is located.
 - In the work pane, select the mobile partition.
 - From the Tasks menu, click **Properties**. The Partition Properties window is displayed.
 - Click the **Virtual Adapters** tab.
 - Record the **Connecting Partition** and the **Connecting Adapter** for each virtual adapter on the mobile partition.
 - The **Connecting Partition** is the Virtual I/O Server logical partition that contains the server virtual adapter to which the virtual adapter on the mobile partition connects.
 - The **Connecting Adapter** is the ID of the virtual adapter on the Virtual I/O Server logical partition to which the virtual adapter on the mobile partition connects.

An example follows:

Table 29. Example information for virtual adapters on the mobile partition

Adapter ID	Connecting Partition	Connecting Adapter
2	VIOS1	11
4	VIOS1	12

- Click **OK** to exit the Partition Properties window.
- Verify the virtual adapter configuration of each Connecting Partition, or Virtual I/O Server logical partition, that you identified in the previous step:
 - In the navigation pane, expand **Systems Management > Servers**.

- b. Click the managed system on which the mobile partition is located.
- c. In the work pane, select a Virtual I/O Server logical partition from which the mobile partition receives virtual I/O resources.
- d. From the Tasks menu, click **Properties**. The Partition Properties window is displayed.
- e. Click the **Virtual Adapters** tab.
- f. Verify that the virtual adapters on the Virtual I/O Server logical partition are connected to the virtual adapters on the mobile partition:
 - The **Adapter ID** of the virtual adapter on the Virtual I/O Server logical partition corresponds to the **Connecting Adapter** that you recorded for the virtual adapter on the mobile partition.
 - The **Connecting Adapter** of the virtual adapter on the Virtual I/O Server logical partition corresponds to the **Adapter ID** that you recorded for the virtual adapter on the mobile partition. The value for virtual SCSI adapters can also be set to **Any Partition Slot**.

An example follows:

Table 30. Example information for virtual adapters on the Virtual I/O Server logical partition

Adapter ID	Connecting Partition	Connecting Adapter
11	Mobile Partition	2
12	Mobile Partition	Any Partition Slot

- g. Click **OK** to exit the Partition Properties window.
3. If all of the virtual SCSI adapters on the Virtual I/O Server logical partition allow access to virtual SCSI adapters of every logical partition (the **Connecting Partition** for every virtual SCSI adapter is set to **Any Partition**), complete one the following steps:
 - Create a new virtual SCSI adapter on the Virtual I/O Server logical partition and allow only a virtual SCSI adapter on the mobile partition to access it.
 - Change the connection specifications of a virtual SCSI adapter on the Virtual I/O Server logical partition so that it only allows access to a virtual SCSI adapter on the mobile partition.

Verifying that the mobile partition has access to its physical storage:

You can use the Hardware Management Console (HMC) to verify that the mobile partition has access to its physical storage on the storage area network (SAN) so that the mobile partition can access its physical storage after it moves to the destination server.

For partition mobility to be successful, the mobile partition must have access to the same physical storage from both the source and destination environments. In the source environment, the following connections must exist:

- Each virtual SCSI adapter on the mobile partition must have access to a target virtual SCSI adapter on the source Virtual I/O Server logical partition.
- The target virtual SCSI adapters on the source Virtual I/O Server logical partition must have access to a SAN host-attached adapter on the source Virtual I/O Server logical partition.
- The SAN host-attached adapter on the source Virtual I/O Server logical partition must be connected to a storage area network and have access to the physical storage devices you want the mobile partition to have access to in the storage area network.

In the destination environment, the following connections must exist:

- The destination Virtual I/O Server logical partition has unused Virtual slots available.
- The SAN host-attached adapter on the destination Virtual I/O Server logical partition must be connected to the same storage area network as the source Virtual I/O Server logical partition and have access to the same mobile partition physical storage as the source Virtual I/O Server logical partition.

You must be a super administrator to complete this task.

To verify these connections using the HMC, complete the following steps:

1. In the navigation pane, open **Systems Management** and select **Servers**.
2. Select the managed server of your choice in the navigation pane.
3. In the work pane, select the source Virtual I/O Server, click the **Tasks** button, and select **Hardware (Information) > Virtual Adapters > SCSI**.
4. Verify the following information and click **OK**:
 - Virtual Adapter
 - Backing Device
 - Remote Partition
 - Remote Adapter
 - Remote Backing Device

Tip: The virtual SCSI adapter fields might be blank if the mobile partition is powered off or if the physical disk has not been linked to the Virtual I/O Server's virtual SCSI adapter.

If the information is incorrect, return to "Preparing the virtual SCSI configuration for partition mobility" on page 74 and complete the task associated with the incorrect information.

Specifying a new name for a virtual target device to use on a destination VIOS partition:

Before you move a logical partition, you can specify a new name for a virtual target device, if needed. After you move the logical partition, the virtual target device assumes the new name on the Virtual I/O Server (VIOS) partition on the destination system.

Before you start, verify that the following products are at the following versions:

- The Hardware Management Console (HMC) is at version 7 release 3.5.0, or later.
- The VIOS partitions are at version 2.1.2.0, or later. This requirement applies to both the source VIOS partitions and the destination VIOS partitions.

Where possible, partition mobility preserves user-defined names of the virtual target devices on the destination system. Partition mobility does not preserve vtscsix IDs.

In some situations, partition mobility might not be able to preserve a user-defined name. For example, when the name is already in use on the destination VIOS partition.

If you want to maintain user-defined names on the destination VIOS partition, you can specify a new name for the virtual target device to use on the destination VIOS partition. If you do not specify a new name, partition mobility automatically assigns the next available vtscsix name to the virtual target device on the destination VIOS partition.

1. To view the names and mappings of the virtual target devices, run the **lsmmap** command as follows. Run the command from the command-line interface on the source VIOS partition:

```
lsmmap -all
```

The output might look like the following output:

SVSA	Physloc	Client Partition ID
vhost4	U8203.E4A.10D4431-V8-C14	0x0000000d
VTD	client3_hd0	
Status	Available	
LUN	0x8100000000000000	
Backing device	hdisk5	
Physloc	U789C.001.DQ1234#-P1-C1-T1-W500507630508C075-L4002402300000000	

```

VTD                client3_hd1
Status             Available
LUN                0x8200000000000000
Backing device     hdisk6
Physloc            U789C.001.DQ1234#-P1-C1-T1-W500507630508C075-L4002402400000000

```

In this example, the user-defined names of the virtual target devices are `client3_hd0` and `client3_hd1`.

2. To specify a user-defined name for a virtual target device to use on the destination VIOS partition, run the **chdev** command as follows. Run the command from the command-line interface on the source VIOS partition:

```
chdev -dev dev_id -attr mig_name=partition_mobility_id
```

where:

- *dev_id* is the user-defined name of the virtual target device on the source VIOS partition.
- *partition_mobility_id* is the user-defined name that you want the virtual target device to have on the destination VIOS partition.

Preparing the virtual Fiber Channel configuration for partition mobility

You need to verify that the virtual Fiber Channel configuration is configured correctly so that you can successfully move the mobile partition from the source server to the destination server by using the Hardware Management Console (HMC).

Before you plan an inactive partition migration on a logical partition that has an N_Port ID Virtualization (NPIV) adapter, you must ensure that the logical partition had been activated at least once.

The verification includes tasks such as verifying the worldwide port names (WWPNs) of the virtual Fiber Channel adapters on the mobile partition, and verifying that the physical Fibre Channel adapters and the physical Fiber Channel switches support NPIV. AIX NPIV clients using AIX Multipath I/O (MPIO) software support single path reserve with live partition mobility.

You can migrate a client partition that has mapped NPIV adapters for which no WWPN targets have been zoned, by specifying the Fiber Channel port to be used on the destination partition. If the physical port which must be used on the destination partition is specified, validation checks the physical port to ensure that it has no WWPN targets that are zoned and that the virtual adapter is mapped on the destination partition. When the physical port is not specified, validation checks all ports on the destination partition to determine whether there are any WWPN targets that are zoned. If any WWPN targets that are zoned are found, the validation fails. If there are no WWPN targets that are zoned, the virtual adapter is not mapped on the destination partition.

The destination server must provide the same virtual Fiber Channel configuration as the source server so that the mobile partition can access its physical storage on the storage area network (SAN) after it moves to the destination server.

To prepare the virtual Fiber Channel configuration for active or inactive partition mobility, complete the following tasks.

Table 31. Preparation tasks for the virtual Fiber Channel configuration on systems that are managed by the HMC


Storage planning tasks	Active mobility task	Inactive mobility task	Information resources
1. For each virtual Fiber Channel adapter on the mobile partition, verify that both (active and inactive) WWPNs for each virtual Fiber Channel adapter is zoned to the same storage port WWN on the switch or the storage area network (SAN).	X	X	<ul style="list-style-type: none"> “Identifying the WWPNs that are assigned to a virtual Fibre Channel adapter” IBM System Storage SAN Volume Controller
2. Verify that the physical Fibre Channel adapters that are assigned to the source and destination Virtual I/O Server logical partitions support NPIV. Run the lsnports command to view the physical ports on the physical Fiber Channel adapters that support NPIV.	X	X	Virtual I/O Server and Integrated Virtualization Manager commands
3. Verify that the switches to which the physical Fibre Channel adapters on both the source and destination Virtual I/O Server logical partitions are cabled support NPIV. Run the lsnports command to view the fabric support of the physical ports on the physical Fiber Channel adapters. If the fabric support is 1, the physical port is cabled to a switch that supports NPIV.	X	X	Virtual I/O Server and Integrated Virtualization Manager commands
4. Verify that the mobile partition has access to the virtual Fiber Channel adapters on the source Virtual I/O Server logical partition.	X	X	“Verifying the virtual adapter connections between the mobile partition and the Virtual I/O Server logical partitions on the source server” on page 76
5. If you changed any partition profile attributes, restart the mobile partition for the new values to take effect.	X	X	Shutting down and restarting logical partitions

Related concepts:

“Storage configuration in a partition mobility environment” on page 43

Learn about the virtual SCSI and virtual Fibre Channel configuration required for partition mobility that is managed by the Hardware Management Console (HMC).

Related information:

 Redundancy configuration using virtual fibre channel adapters

Identifying the WWPNs that are assigned to a virtual Fibre Channel adapter:

You can identify the worldwide port names (WWPNs) that are assigned to the virtual Fibre Channel adapters on the mobile partition by using the Hardware Management Console (HMC) to view the partition properties of the mobile partition.

To identify the WWPNs that are assigned to a virtual Fibre Channel adapter using the HMC, complete the following steps:

1. In the navigation pane, expand **Systems Management > Servers**.
2. Click the server on which the mobile partition is located.
3. In the navigation pane, select the mobile partition.
4. From the **Tasks** menu, click **Properties**. The Partition Properties window is displayed.
5. Click the **Virtual Adapters** tab.
6. Select a virtual Fibre Channel adapter.
7. From the **Actions** menu, click **Properties**. The Virtual Fibre Channel Adapter Properties window is displayed.

8. Repeat steps 6 on page 80 and 7 on page 80 for each virtual Fibre Channel adapter on the mobile partition.
9. Click **Close** to return to the Partition Properties window.

Verifying the virtual adapter connections between the mobile partition and the Virtual I/O Server logical partitions on the source server:

You can verify the virtual adapter connections between the mobile partition and the Virtual I/O Server logical partitions on the source server so that the Hardware Management Console (HMC) correctly configures the virtual adapters on the destination server when you move the mobile partition.

To verify the virtual adapter connections between the mobile partition and the source Virtual I/O Server logical partitions, complete the following steps from the HMC:

1. Verify the virtual adapter configuration of the mobile partition:
 - a. In the navigation pane, expand **Systems Management > Servers**.
 - b. Click the managed system on which the mobile partition is located.
 - c. In the work pane, select the mobile partition.
 - d. From the Tasks menu, click **Properties**. The Partition Properties window is displayed.
 - e. Click the **Virtual Adapters** tab.
 - f. Record the **Connecting Partition** and the **Connecting Adapter** for each virtual adapter on the mobile partition.
 - The **Connecting Partition** is the Virtual I/O Server logical partition that contains the server virtual adapter to which the virtual adapter on the mobile partition connects.
 - The **Connecting Adapter** is the ID of the virtual adapter on the Virtual I/O Server logical partition to which the virtual adapter on the mobile partition connects.

An example follows:

Table 32. Example information for virtual adapters on the mobile partition

Adapter ID	Connecting Partition	Connecting Adapter
2	VIOS1	11
4	VIOS1	12

- g. Click **OK** to exit the Partition Properties window.
2. Verify the virtual adapter configuration of each Connecting Partition, or Virtual I/O Server logical partition, that you identified in the previous step:
 - a. In the navigation pane, expand **Systems Management > Servers**.
 - b. Click the managed system on which the mobile partition is located.
 - c. In the work pane, select a Virtual I/O Server logical partition from which the mobile partition receives virtual I/O resources.
 - d. From the Tasks menu, click **Properties**. The Partition Properties window is displayed.
 - e. Click the **Virtual Adapters** tab.
 - f. Verify that the virtual adapters on the Virtual I/O Server logical partition are connected to the virtual adapters on the mobile partition:
 - The **Adapter ID** of the virtual adapter on the Virtual I/O Server logical partition corresponds to the **Connecting Adapter** that you recorded for the virtual adapter on the mobile partition.
 - The **Connecting Adapter** of the virtual adapter on the Virtual I/O Server logical partition corresponds to the **Adapter ID** that you recorded for the virtual adapter on the mobile partition. The value for virtual SCSI adapters can also be set to **Any Partition Slot**.

An example follows:

Table 33. Example information for virtual adapters on the Virtual I/O Server logical partition

Adapter ID	Connecting Partition	Connecting Adapter
11	Mobile Partition	2
12	Mobile Partition	Any Partition Slot

- g. Click **OK** to exit the Partition Properties window.
3. If all of the virtual SCSI adapters on the Virtual I/O Server logical partition allow access to virtual SCSI adapters of every logical partition (the **Connecting Partition** for every virtual SCSI adapter is set to **Any Partition**), complete one the following steps:
 - Create a new virtual SCSI adapter on the Virtual I/O Server logical partition and allow only a virtual SCSI adapter on the mobile partition to access it.
 - Change the connection specifications of a virtual SCSI adapter on the Virtual I/O Server logical partition so that it only allows access to a virtual SCSI adapter on the mobile partition.

Validating the configuration for partition mobility

You can use the Partition Migration wizard on the Hardware Management Console (HMC) to validate the configuration of the source and destination systems for partition mobility. If the HMC detects a configuration or connection problem, it displays an error message with information to help you resolve the problem.

If the source and destination servers are managed by different HMCs, verify that the Secure Shell (SSH) authentication keys are set up correctly between the HMCs. For instructions, see “Verifying SSH authentication between the source and destination HMCs” on page 58.

You must be a super administrator to validate the partition mobility environment.

To validate the source and destination systems for partition mobility using the HMC, complete the following steps:

1. In the navigation pane, open **Systems Management**.
2. Select **Servers**.
3. In the navigation pane, select the source server.
4. Select the mobile partition and expand **Operations > Mobility > Validate**. The Partition Migration Validation window opens.
5. Specify information about the partition mobility environment, and then click **Validate**. The Virtual Storage assignments table populates with suggested virtual adapter settings.

Remember: With HMC version 7 release 3.5.0, or later, you can select **Override virtual storage errors when possible**. Select this option to validate moving the mobile partition to a destination system with less redundancy.

6. Review the available virtual adapter settings on the destination system.
7. Click **Validate** again to confirm that the changed settings are still acceptable for partition mobility.

Where possible, the HMC Version 7 Release 3.5.0, or later, preserves the virtual slot assignments of the virtual server adapters on the destination system. However, in some situations the HMC might not be able to preserve a virtual slot ID. For example, when the slot ID is already occupied on the destination Virtual I/O Server (VIOS) logical partition. When the HMC cannot preserve a virtual slot ID, you receive an error message, and the HMC assigns an available slot ID. You can override the assignments by completing the following steps from the HMC command-line interface:

1. Run the `lslparmigr` command to show a list of available slot IDs for a VIOS partition.
2. Run the `migr1par` command to accomplish the following tasks:
 - Specify virtual slot IDs for one or more virtual adapter mappings.

- Validate the specified slot IDs.

Note: You can specify the port name of the Fibre Channel to be used for creating Fibre Channel mapping on the destination server when you are performing partition migration.

You can use the HMC command-line interface to specify the port name.

- List all the valid port names of the Fibre Channel by running the **lsnports** command.
- From the list of valid port names, specify the port name you want to use in the **vios_fc_port_name** attribute, by running the following command:

```
migr1par -o v -m <srcCecName> -t <dstCecName> -p <lparName> -i "virtual_fc_mappings=
<Client_slot_num>/<target_vios_name>/<target_vios_id>/<target_slot_num>/<vios_fc_port_name>"
```

For example:

```
migr1par -o v -m vrml13-fsp -t vrml11-fsp -p vrml11p03 -i "virtual_fc_mappings=
3/vrml11-vios1/1/8/fcs0"
```

Related concepts:

“Configuration validation for partition mobility” on page 6

You can learn about the tasks that the Partition Migration wizard on the Hardware Management Console (HMC) performs to validate your system configuration for active and inactive partition mobility.

Related tasks:

“Specifying a new name for a virtual target device to use on a destination VIOS partition” on page 78

Before you move a logical partition, you can specify a new name for a virtual target device, if needed.

After you move the logical partition, the virtual target device assumes the new name on the Virtual I/O Server (VIOS) partition on the destination system.

Moving the mobile partition

You can move an active, inactive, or suspended logical partition from one server to another server by using the Partition Migration wizard on the Hardware Management Console (HMC).

Moving the mobile partition with HMC

You can move an active or inactive logical partition from one server to another server by using the Partition Migration wizard on the Hardware Management Console (HMC).

Before you move a logical partition from one server to another server, complete the following tasks from the HMC.

Table 34. Prerequisite tasks for moving a logical partition

Partition mobility prerequisite tasks	Active mobility task	Inactive mobility task	Information resources
1. Verify that you have completed all the required preparation tasks for partition mobility.	X	X	“Preparing for partition mobility” on page 48
2. Verify that the source and destination servers are in the Operating state.	X	X	To power on a managed system, see Power on
3. Verify that the mobile partition is powered off. Requirement: Return the logical partition to an Operating state when the following conditions are true: <ul style="list-style-type: none"> • You want to actively move the logical partition. • The logical partition is in a failed state 		X	<ul style="list-style-type: none"> • Shutting down and restarting logical partitions • Reference Code finder
4. Verify that the mobile partition is in the Operating state.	X		Activating a logical partition using the HMC

Table 34. Prerequisite tasks for moving a logical partition (continued)

Partition mobility prerequisite tasks	Active mobility task	Inactive mobility task	Information resources
5. Verify that the source and destination Virtual I/O Servers are active.	X	X	Activating a logical partition using the HMC
6. Verify that all tape and CD jobs are completed or stopped.	X		
7. Verify that no dynamic logical partitioning (DLPAR) operations are running on any of the logical partitions on the source server and destination server. Do not perform any DLPAR operations on any of the logical partitions on the source server and the destination server during partition mobility. You can perform DLPAR operations on the logical partitions after the mobile partition successfully moves to the destination server.	X	X	
8. If the source and the destination servers are managed by different HMCs, verify that the Secure Shell (SSH) authentication keys are set up correctly between the HMCs.	X	X	“Verifying SSH authentication between the source and destination HMCs” on page 58
9. Run the migration verification tool on the HMC to verify that the servers, Virtual I/O Servers, mobile partition, storage, and network are ready for partition mobility.	X	X	“Validating the configuration for partition mobility” on page 82

To move a logical partition from one server to another server by using the HMC, complete the following tasks:

1. In the navigation pane, open **Systems Management**.
2. Select **Servers**.
3. In the work pane, open the source server.
4. Select the mobile partition and select **Operations > Mobility > Migrate**.
5. Complete the wizard.

Tips:

- a. With HMC version 7 release 3.5.0, or later, you can select **Override virtual storage errors when possible**. Select this option if you want to move the mobile partition to a destination system with less redundancy.
- b. Where possible, the HMC version 7 release 3.5.0, or later, preserves the virtual slot assignments of the virtual server adapters on the destination system. However, in some situations the HMC might not be able to preserve one or more virtual slot IDs. In this situation, the HMC assigns available slot IDs. To override the assignments, move the mobile partition by running the **migr1par** command from the HMC command-line interface.
- c. You can specify the IP address of the mover service partition on the source server, the mover service partition on the destination server, or both. For example, you want partition mobility to use the fastest IP address available on a mover service partition. To specify the IP address of a mover service partition, the following products must be at the specified versions:
 - The HMC must be at version 7 release 3.5.0, or later.
 - The mover service partition for which you specify an IP address must be at Virtual I/O Server version 2.1.2.0, or later.

To specify the IP addresses of the mover service partitions, move the mobile partition by running the **migr1par** command from the HMC command-line interface.

After you move a logical partition from one server to another server, complete the following tasks.

Table 35. Postrequisite tasks for moving a logical partition

Partition mobility postrequisite tasks	Active mobility task	Inactive mobility task	Information resources
1. Activate the mobile partition on the destination server.		X	Activating a logical partition using the HMC
2. Optional: Add dedicated I/O adapters and single root I/O virtualization (SR-IOV) logical ports to the mobile partition on the destination server.	X	X	<ul style="list-style-type: none"> • Adding physical I/O devices and slots dynamically • Adding a single root I/O virtualization logical port to a logical partition dynamically
3. If any virtual terminal connections were lost during the migration, re-establish the connections on the destination server.	X	X	
4. Optional: Assign the mobile partition to a logical partition workload group.	X	X	“Adding the mobile partition to a partition workload group”
5. If mobility-unaware applications terminated on the mobile partition before its movement, restart those applications on the destination.	X		
6. If you changed any partition profile attributes, shut down and activate the new profile for the new values to take effect.	X	X	Shutting down and restarting logical partitions
7. Optional: Back up the Virtual I/O Server logical partitions on the destination server to preserve the new virtual device mappings.	X	X	Backing up the Virtual I/O Server
8. Optional: Disable secure IP tunnels between the mover service partitions on the source and destination servers.	X		stopsvc command

Adding the mobile partition to a partition workload group:

You can add the mobile partition to a partition workload group by using the Hardware Management Console (HMC) after you move the mobile partition from the source server to the destination server.

You must be a super administrator to complete this task.

A partition workload group identifies a set of logical partitions that are located on the same physical system. Workload management tools use partition workload groups to identify which logical partitions they can manage.

Prior to moving the mobile partition from the source environment to the destination environment, you might have removed the mobile partition from a partition workload group. Now that you have successfully moved the mobile partition to the destination environment, you can add it to a partition workload group.

To add the mobile partition to a partition workload group using the HMC, complete the following steps:

1. In the navigation pane, open **Systems Management** and select **Servers**.
2. Select the managed server of your choice in the navigation pane.

3. Select the logical partition of your choice in the work pane.
4. Select **Configuration > Manage Profiles**.
5. Select the profile of your choice and select **Actions > Edit**.
6. Click the **Settings** tab.
7. In the Workload Management area, select (None) and click **OK**.
8. Repeat steps 1 through 7 for all partition profiles associated with the mobile partition. In order for this change to take effect, you will need to activate this logical partition with this profile.

This can also be changed using DLPAR by selecting the logical partition > **Properties > Other** tab.

Moving the suspended mobile partition with the HMC command-line interface

You can move a suspended Linux logical partition from one server to another server by using the Hardware Management Console (HMC) command-line interface.

Note: Moving a suspended logical partition to another managed system exposes the logical partition to accidental reassignment of its virtual storage devices while it remains suspended. Because this exposure cannot be prevented, it is preferred that the suspended logical partition be resumed before the logical partition is moved.

You can suspend a Linux logical partition with its operating system and applications, and store its virtual server state to persistent storage. At a later stage, you can resume the operation of the logical partition.

To move a suspended logical partition from one managed system to the other, you can run the **migr1par** command with the `protectstorage` attribute set to a value of 2. Since the virtual storage devices assigned to the suspended logical partition are no longer protected after the suspended logical partition has been moved, you must ensure the integrity of the virtual storage devices while the logical partition remains suspended.

After you move a suspended logical partition from one server to another server, you can perform one of the following actions:

- Resume the mobile partition on the destination server. For instructions, see “Resuming the suspended mobile partition with HMC.”
- Shut down the mobile partition on the destination server. For instructions, see “Shutting down the suspended mobile partition with HMC” on page 87.

Related tasks:

“Resuming the suspended mobile partition with HMC”

You can resume a suspended Linux logical partition on the server by using the Hardware Management Console (HMC) Version 7.7.2.0, or later.

“Shutting down the suspended mobile partition with HMC” on page 87

You can shut down a suspended Linux logical partition on the server by using the Hardware Management Console (HMC) Version 7.7.2.0, or later.

Resuming the suspended mobile partition with HMC:

You can resume a suspended Linux logical partition on the server by using the Hardware Management Console (HMC) Version 7.7.2.0, or later.

To resume a suspended logical partition on the server by using the HMC, complete the following tasks:

1. In the navigation pane, open **Systems Management**.
2. Select **Servers**.
3. In the work pane, select the suspended mobile partition.
4. Select **Operations > Suspend Operations > Resume**.


Note: If the Virtual Station Interface (VSI) configuration on the destination server fails, the resume operation also fails. You must then shut down and restart the partition in order to recover from the failed resume operation.

Related tasks:

“Moving the suspended mobile partition with the HMC command-line interface” on page 86
You can move a suspended Linux logical partition from one server to another server by using the Hardware Management Console (HMC) command-line interface.

“Determining the virtual Ethernet switch name and mode in the destination server” on page 128
Determine the name and mode of the virtual Ethernet switches in the destination server by using the Hardware Management Console (HMC).

Related information:

 Suspending a logical partition

Shutting down the suspended mobile partition with HMC:

You can shut down a suspended Linux logical partition on the server by using the Hardware Management Console (HMC) Version 7.7.2.0, or later.

To shut down a suspended logical partition on the server by using the HMC, complete the following tasks:

1. In the navigation pane, open **Systems Management**.
2. Select **Servers**.
3. In the work pane, select the suspended mobile partition.
4. Select **Operations > Shut Down**.

Related tasks:

“Moving the suspended mobile partition with the HMC command-line interface” on page 86
You can move a suspended Linux logical partition from one server to another server by using the Hardware Management Console (HMC) command-line interface.

Troubleshooting partition mobility

Learn how to understand, isolate, and resolve problems related to active and inactive partition mobility by using the Hardware Management Console (HMC).

Sometimes you will be able to resolve a problem on your own, while at other times you will need to gather information to help the service technicians resolve your problem in a timely manner.

Troubleshooting active partition mobility

Learn how to troubleshoot problems that might occur with active partition mobility by using the Hardware Management Console (HMC).

The following table lists possible errors and ways to recover.

Table 36. Known problems and solutions for active partition mobility

Problem	Solution
<p>You receive the following error:</p> <p>HSC3656 There is an insufficient amount of memory available on the destination managed system for the configuration of the partition. Please perform one or both of the following actions: 1. Remove memory from any shutdown dedicated memory partitions on the destination managed system. 2. Remove memory from any running dedicated memory partitions on the destination managed system.</p>	<ol style="list-style-type: none"> 1. To make physical memory available for the mobile partition, dynamically remove physical memory from inactive logical partitions that use dedicated memory (subsequently referred to as <i>dedicated memory partitions</i>) on the destination server by running the chhwres command from the HMC command line. For example, <code>chhwres -r mem -m <destination_server> -o r -p <logical_partition> -q <memory></code>, where: <ul style="list-style-type: none"> • <i><destination_server></i> is the name of the server to which you want to move the mobile partition. • <i><logical_partition></i> is the name of the logical partition from which you want to remove physical memory. • <i><memory></i> is the amount of physical memory, in MB, that you want to remove from the logical partition. 2. If you cannot satisfy the memory requirement of the mobile partition by removing physical memory from dedicated memory partitions that are inactive, dynamically remove physical memory from dedicated memory partitions that are active on the destination server by performing one of the following tasks: <ul style="list-style-type: none"> • Removing dedicated memory dynamically using the HMC • Running the chhwres command from the HMC command line.

Table 36. Known problems and solutions for active partition mobility (continued)

Problem	Solution
<p>You receive the following error:</p> <pre>HSCL03EC There is not enough memory: Obtained : xxxx, Required : xxxx. Check that there is enough memory available to activate the partition. If not, create a new profile or modify the existing profile with the available resources, then activate the partition. If the partition must be activated with these resources, de-activate any running partition(s) using the resource then activate this partition.</pre>	<ol style="list-style-type: none"> 1. To make physical memory available for the mobile partition, dynamically remove physical memory from inactive logical partitions that use dedicated memory (subsequently referred to as <i>dedicated memory partitions</i>) on the destination server by running the chhwres command from the HMC command line. For example, <code>chhwres -r mem -m <destination_server> -o r -p <logical_partition> -q <memory></code>, where: <ul style="list-style-type: none"> • <i><destination_server></i> is the name of the server to which you want to move the mobile partition. • <i><logical_partition></i> is the name of the logical partition from which you want to remove physical memory. • <i><memory></i> is the amount of physical memory, in MB, that you want to remove from the logical partition. 2. If you cannot satisfy the memory requirement of the mobile partition by removing physical memory from dedicated memory partitions that are inactive, dynamically remove physical memory from dedicated memory partitions that are active on the destination server by performing one of the following tasks: <ul style="list-style-type: none"> • Removing dedicated memory dynamically using the HMC • Running the chhwres command from the HMC command line. 3. If you cannot satisfy the memory requirement of the mobile partition by dynamically removing physical memory from dedicated memory partitions that are active on the destination server, dynamically remove memory from the mobile partition. For instructions, see Removing dedicated memory dynamically using the HMC. 4. If you cannot reduce the amount of memory required by the mobile partition to an amount that is equal to or less than the amount of memory that is available on the destination server, shut down logical partitions on the destination server until enough memory is available for the mobile partition to activate on the destination server. 5. If you cannot satisfy the memory requirement of the mobile partition by shutting down logical partitions on the destination server, move the mobile partition to the destination server by using inactive partition mobility.

Table 36. Known problems and solutions for active partition mobility (continued)

Problem	Solution
	<p>Notes:</p> <ol style="list-style-type: none"> 1. The mobile partition must use dedicated memory. If the mobile partition uses shared memory, skip step 3 on page 89 and continue to the next step. 2. After you move the logical partition to the destination server, you might be able to dynamically add one logical memory block (LMB) back to the logical partition. This can happen in one or more of the following situations: <ul style="list-style-type: none"> • The actual available LMBs on the destination server are fractionally high. When determining the available LMBs on the destination server, all fractional LMB sizes are rounded down to the nearest whole number. For example, 5.9 LMBs are rounded down to 5 LMBs. • The amount of internal hypervisor storage used on the destination server (to support the logical partition) is a small fraction of 1 LMB. When determining the amount of memory required by the logical partition on the destination server, one LMB is added to the actual number of LMBs required by the logical partition. The added LMB accounts for internal hypervisor storage required to support the logical partition on the destination server.


Table 36. Known problems and solutions for active partition mobility (continued)

Problem	Solution
<p>You receive the following error:</p> <p>HSCLA319 The migrating partition's virtual Fibre Channel client adapter cannot be hosted by the existing Virtual I/O Server (VIOS) partitions on the destination managed system.</p>	<p>This error indicates that the VIOSs in the target server do not have suitable resources to host the virtual Fibre Channel adapter on the migrating or suspended partition. The following are the most common reasons for this error:</p> <ul style="list-style-type: none"> • The storage area network (SAN) employs port zoning. The target server ports and source server ports are not zoned identically. To host the migrating virtual adapter, the list of Fibre Channel targets in a port on the target server must exactly match the list of Fibre Channel targets in the current mapped port of the migrating virtual adapter on the source server. • The two worldwide port names (WWPNs) assigned to the virtual adapter are not zoned identically. The two WWPNs must be interchangeable from both SAN and storage point of view. • The target server does not have a port that can meet or exceed the maximum transfer size of the source server port. The maximum transfer size is an attribute of the Fibre Channel port and can be viewed by running the lsattr command on a Fibre Channel device. • A switch on the SAN might be configured to use features that extend the Fibre Channel standard in ways that are not compatible with Live Partition Mobility. For example, a port binding feature that tracks WWPN-to-port mappings. This feature can cause problems because Live Partition Mobility validation requires that all ports must be explored through a series of login and logout operations. If the switch tries to track the WWPN-to-port mappings, it might run out of resources and not permit login operations. Disabling this type of feature solves some problems related to failed Fibre Channel login operations.
<p>If the operating system running in the mobile partition does not explicitly support the processor version register of the destination server, and the processor determines that explicit support is required, the processor will not allow the migration to proceed.</p>	<p>Perform one of the following actions:</p> <ul style="list-style-type: none"> • Move the logical partition to another system. • Update the operating system to a level that supports the target system processor version registers.
<p>You receive an error pertaining to the operating system when you attempt to migrate a logical partition.</p>	<ol style="list-style-type: none"> 1. Examine the operating system error logs for operating system-related failures. 2. Examine the HMC log for application-related failures.
<p>You receive an HMC error pertaining to insufficient physical memory on the destination server. Important: Sufficient physical memory includes the amount of available physical memory on the server and the amount of available contiguous physical memory on the server. If the mobile partition requires more contiguous physical memory, making more physical memory available will not solve the problem.</p>	<p>Perform one of the following actions:</p> <ul style="list-style-type: none"> • Move the logical partition to a different server. • Make more physical memory available on the destination server. See “Determining the available physical memory on the destination server” on page 53 for instructions.

Table 36. Known problems and solutions for active partition mobility (continued)

Problem	Solution
<p>The HMC (or HMCs) and managed system lost their connection while the migration was in progress, or the migration failed.</p>	<p>Before running migration recovery ensure that the Resource Monitoring and Control (RMC) connections are established for the migrating partition and the VIOS partitions on the source and destination servers. Complete the following steps from the HMC that manages the source server. If the source server or the source HMC are unavailable, complete the following steps from the HMC that manages the destination server.</p> <ol style="list-style-type: none"> 1. In the navigation pane, open Systems Management. 2. Select Servers. 3. In the work pane, select the source server. If the source server is unavailable, select the destination server. 4. In the Tasks menu, select Mobility > Recover. The Migration Recovery window is displayed. 5. Click Recover. 6. If you recovered the migration from the HMC that manages the destination server (and a different HMC manages the source server), you might have to manually perform additional recovery tasks on the source server to finish the recovery. For example, even though the migration occurs and the mobile partition runs on the destination server, the mobile partition might appear as an inactive logical partition on the source server. In this situation, remove the mobile partition from the source server to finish the recovery. <p>Tip: You can also run the <code>migr1par -o r</code> command to recover a migration.</p> <p>Note: When you migrate a partition remotely, ensure that you do not connect the source and target servers to the same HMC.</p>
<p>While attempting to change resources dynamically, you receive an error that the RMC daemon is not connected.</p>	<p>This error typically occurs when there is a network connection problem between the logical partitions and the HMC. To resolve this error, check your system network setup.</p>
<p>Live Partition Mobility fails when the client logical partition has multiple virtual Fibre Channel adapters mapped to the same physical Fibre Channel adapter.</p>	<p>You cannot migrate or suspend logical partitions that have multiple virtual Fibre Channel adapters mapped to the same physical Fibre Channel adapter.</p>

Related information:

 [Partition mobility firmware support matrix](#)

Troubleshooting inactive partition mobility

Learn how to troubleshoot problems with inactive partition mobility using the Hardware Management Console (HMC).

The following table lists possible errors and ways to recover.

Table 37. Known problems and solutions for inactive partition mobility

Problem	Solution
If the mobile partition is moved to a server that the operating system does not support (and explicit support is required), then the boot of the logical partition on the destination server will fail.	Move the logical partition to a different server.
You receive an HMC error pertaining to insufficient physical memory on the destination server. Important: Sufficient physical memory includes the amount of available physical memory on the server and the amount of available contiguous physical memory on the server. If the mobile partition requires more contiguous physical memory, making more physical memory available will not solve the problem.	Perform one of the following actions: <ul style="list-style-type: none"> • Move the logical partition to a different server. • Make more physical memory available on the destination server. See “Determining the available physical memory on the destination server” on page 53 for instructions.

Virtual I/O Server errors

Learn about the errors that might occur on the Virtual I/O Server (VIOS).

The following table lists possible VIOS errors and their definitions.

Table 38. VIOS error codes

Error Code	Definition
1	The virtual adapter is not ready to be moved. The source virtual Ethernet is not bridged.
2	The virtual adapter can be moved with less capability. All virtual local area networks (VLAN) are not bridged on the destination. Hence, the virtual Ethernet adapter has less capability on the target system compared to the source system.
3	The stream ID is still in use.
64	The migmgr command cannot be started.
65	The stream ID is invalid.
66	The virtual adapter type is invalid.
67	The virtual adapter DLPAR resource connector (DRC) name is not recognized.
68	The virtual adapter method cannot be started, or it was prematurely terminated.
69	There is a lack of resources (that is, the ENOMEM error code).
80	The storage that is being used by the adapter is specific to the VIOS and cannot be accessed by another VIOS. Hence, the virtual adapter cannot complete the mobility operation.
81	The virtual adapter is not configured.
82	The virtual adapter cannot be placed in a migration state.
83	The virtual devices are not found.
84	The virtual adapter VIOS level is insufficient.
85	The virtual adapter cannot be configured.
86	The virtual adapter is busy and cannot be unconfigured.
87	The virtual adapter or device minimum patch level is insufficient.
88	The device description is invalid.
89	The command argument is invalid.
90	The virtual target device cannot be created because of incompatible backing device attributes. Typically, this is because of a mismatch in the maximum transfer (MTU) size or SCSI reserve attributes of the backing device between the source VIOS and the target VIOS.
91	The DRC name passed to the migration code is for an adapter that exists.

Live Partition Mobility on IVM-managed systems

You can use the Integrated Virtualization Manager (IVM) to move an active or inactive logical partition from one server to another.

Partition mobility overview for IVM

You can learn about the benefits of partition mobility, how the Integrated Virtualization Manager (IVM) performs active and inactive partition mobility, and about the configuration that is required to successfully move a logical partition from one system to another.

Related tasks:

“Preparing for partition mobility” on page 123

You need to verify that the source and destination systems are configured correctly so that you can successfully move the mobile partition from the source system to the destination system. This includes verifying the configuration of the source and destination servers, the Integrated Virtualization Manager (IVM) management partitions, the mobile partition, the virtual storage configuration, and the virtual network configuration.

Benefits of partition mobility

Partition mobility provides systems management flexibility and is designed to improve system availability.

For example:

- You can avoid planned outages for hardware or firmware maintenance by moving logical partitions to another server and then performing the maintenance. Partition mobility can help because you can use it to work around scheduled maintenance activities.
- You can avoid downtime for a server upgrade by moving logical partitions to another server and then performing the upgrade. This allows you to continue your work without disruption.
- If a server indicates a potential failure, you can move its logical partitions to another server before the failure occurs. Partition mobility can help avoid unplanned downtime.
- You can consolidate workloads running on several small, under used servers onto a single large server.
- You can move workloads from server to server to optimize resource use and workload performance within your computing environment. With active partition mobility, you can manage workloads with minimal downtime.
- For some systems, you can move applications from one server to an upgraded server by using PowerVM Live Partition Mobility, without affecting availability of the applications.

However, while partition mobility provides many benefits, it does not do the following functions:

- Partition mobility does not provide automatic workload balancing.
- Partition mobility does not provide a bridge to new functions. Logical partitions must be restarted and possibly reinstalled to take advantage of new features.

Partition mobility process for IVM

Learn about how the Integrated Virtualization Manager (IVM) moves an active or inactive logical partition from one server to another server.

The following table describes the steps that take place during the process of active and inactive partition mobility on the IVM.

Table 39. The steps involved in the process of active and inactive partition mobility on the IVM

Partition mobility step	Active mobility step	Inactive mobility step
1. You ensure that all requirements are satisfied and all preparation tasks are completed.	X	X
2. You shut down the mobile partition.		X
3. You initiate partition mobility by starting the migration task on the IVM.	X	X
4. The IVM extracts the physical device description for each physical adapter on the Virtual I/O Server management partition on the source server. The IVM uses the extracted information to determine whether the Virtual I/O Server management partition on the destination server can provide the mobile partition with the same virtual SCSI, virtual Ethernet, and virtual Fibre Channel configuration that exists on the source server. This includes verifying that the Virtual I/O Server management partition on the destination server has enough available slots to accommodate the virtual adapter configuration of the mobile partition. The IVM uses all this information to generate a list of recommended virtual adapter mappings for the mobile partition on the destination server. Where possible, the IVM preserves the following configurations: <ul style="list-style-type: none"> • User-defined names of the virtual target devices. Partition mobility does not preserve vtscsix IDs. • User-defined adapter IDs for virtual server adapters. 	X	X
5. The IVM prepares the source and destination environments for partition mobility. This includes using the virtual adapter mappings from step 4 for mapping the virtual adapters on the mobile partition to the virtual adapters on the Virtual I/O Server management partition on the destination server.	X	X
6. The IVM transfers the logical partition state from the source environment to the destination environment.	In active partition mobility, the following additional steps occur: <ul style="list-style-type: none"> • The source mover service partition extracts the logical partition state information from the source server and sends it to the destination mover service partition over the network. • The destination mover service partition receives the logical partition state information and installs it on the destination server. 	X
7. The IVM suspends the mobile partition on the source server. The source mover service partition continues to transfer the logical partition state information to the destination mover service partition.	X	
8. The hypervisor resumes the mobile partition on the destination server.	X	

Table 39. The steps involved in the process of active and inactive partition mobility on the IVM (continued)

Partition mobility step	Active mobility step	Inactive mobility step
9. The IVM completes the migration. All resources that were consumed by the mobile partition on the source server are reclaimed by the source server: <ul style="list-style-type: none"> The IVM removes the virtual SCSI adapters and the virtual Fibre Channel adapters (that were connected to the mobile partition) from the source Virtual I/O Server management partition. For a mobile partition that uses shared memory, the IVM deactivates the paging space device that was used by the mobile partition and removes the paging space device (if it was automatically created). 	X	X
10. You activate the mobile partition on the destination server.		X
11. You perform postrequisite tasks, such as adding dedicated I/O adapters to the mobile partition or adding the mobile partition to a partition workload group.	X	X

Configuration validation for partition mobility

You can learn about the tasks that the Integrated Virtualization Manager (IVM) performs to validate your system configuration for active and inactive partition mobility.

Before you attempt to migrate an active logical partition, you must validate your environment. You can use the validation function on the IVM to validate your system configuration. If the IVM detects a configuration or connection problem, it displays an error message with information to help you resolve the problem.

The following tables list validation tasks that the IVM performs to verify that the source and destination systems are ready for active or inactive partition mobility.

General compatibility

Table 40. Validation tasks performed by the IVM to verify general compatibility for active and inactive partition mobility

Validation task	Active mobility task	Inactive mobility task
Checks that the resource monitoring and control (RMC) connections are established.	Checks the RMC connections to the mobile partition, the source and destination Virtual I/O Server (VIOS) management partitions, and the connection between the source and destination mover service partitions.	Checks the RMC connections to the source and destination VIOS management partitions.
Checks mobility capability and compatibility.	Checks the source and destination servers, hypervisor, VIOS management partitions, and mover service partitions.	Checks the VIOS management partitions and the hypervisor.
Checks the number of current migrations against the number of supported migrations.	Checks the number of current active migrations against the number of supported active migrations.	Checks the number of current inactive migrations against the number of supported inactive migrations.

Server compatibility

Table 41. Validation tasks performed by the IVM to verify server compatibility for active and inactive partition mobility

Validation task	Active mobility task	Inactive mobility task
Checks that the necessary processing resources are available to create a shell logical partition on the destination system.	X	X
Checks that the necessary memory resources are available to create a shell logical partition on the destination system.	<ul style="list-style-type: none"> For a mobile partition that uses dedicated memory, checks that enough physical memory is available on the destination system. For a mobile partition that uses shared memory, checks that a shared memory pool is configured on the destination server and that it has enough physical memory to satisfy the entitled memory requirements of the mobile partition. 	For a mobile partition that uses dedicated memory, checks that enough physical memory is available on the destination system.
Checks that the necessary I/O adapter resources are available to create a shell logical partition on the destination system. During validation, the IVM extracts the device description for each virtual adapter on the VIOS management partition on the source server. The IVM uses the extracted information to determine whether the VIOS management partition on the destination server can provide the mobile partition with the same virtual SCSI, virtual Ethernet, and virtual Fibre Channel configuration that exists on the source server. This includes verifying that the VIOS management partition on the destination server has enough available slots to accommodate the virtual adapter configuration of the mobile partition.	X	X
Checks that the logical memory block size is the same on the source and destination servers.	X	

Virtual I/O Server compatibility

Table 42. Validation tasks performed by the IVM to verify the source and destination VIOS management partitions for active and inactive partition mobility

Validation task	Active mobility task	Inactive mobility task
Checks that all required I/O devices are connected to the mobile partition through the VIOS management partition. That is, no physical adapters are assigned to the mobile partition and no virtual serial adapters are in virtual slots higher than 1.	X	X
Checks that no virtual SCSI disks are backed by logical volumes and that no virtual SCSI disks are attached to internal disks (not on the SAN).	X	X

Table 42. Validation tasks performed by the IVM to verify the source and destination VIOS management partitions for active and inactive partition mobility (continued)

Validation task	Active mobility task	Inactive mobility task
Checks that the virtual SCSI disks assigned to the logical partition are accessible by the VIOS management partition on the destination server.		X
Checks that the reservation policies of the physical volumes are the same for the source and destination VIOS partitions.	X	X
Checks that the required virtual LAN IDs are available on the destination VIOS management partition.	X	X
Checks that the user-defined names of the virtual target devices on the source VIOS partition can be maintained on the destination VIOS partition.	X	X
Checks that the user-defined adapter IDs of the virtual server adapters on the source VIOS partition can be maintained on the destination VIOS partition.	X	X
For a mobile partition that uses shared memory, the IVM checks for an available paging space device in one of the following ways: <ul style="list-style-type: none"> • Checks that the paging storage pool on the destination server has enough available space to create a paging space device for the mobile partition. • Checks that the management partition on the destination server has access to an available paging space device that meets the size requirements of the mobile partition. 	X	

Mobile partition compatibility

Table 43. Validation tasks performed by the IVM to verify that the mobile partition can successfully move to the destination server by using active or inactive partition mobility

Validation task	Active mobility task	Inactive mobility task
Checks that the operating system on the mobile partition is the Linux operating system.	X	X
Checks the mobile partition, its operating system, and its applications for migration capability.	X	
Checks that the mobile partition is not the redundant error path reporting logical partition.	X	X
Checks that the mobile partition is not in a partition workload group.	X	X
Checks the uniqueness of the virtual MAC addresses or the mobile partition.	X	X
Checks the state of the mobile partition.	Checks that the mobile partition state is Active or Running.	Checks that the mobile partition state is Not Activated.
Checks that the name of the mobile partition is not already in use on the destination server.	X	X
Checks that the mobile partition is not configured with barrier synchronization register (BSR) arrays.	X	

Table 43. Validation tasks performed by the IVM to verify that the mobile partition can successfully move to the destination server by using active or inactive partition mobility (continued)

Validation task	Active mobility task	Inactive mobility task
Checks that the mobile partition is not configured with huge pages.	X	
Checks that the mobile partition does not have a Host Ethernet Adapter (or Integrated Virtual Ethernet).	X	
Checks whether the mobile partition has any connected tape or optical devices as migration fails if any of these devices are connected.	X	X

Related tasks:

“Validating the configuration for partition mobility” on page 142

You can use the Integrated Virtualization Manager (IVM) to validate the configuration of the source and destination systems for partition mobility. If the IVM detects a configuration or connection problem, it displays an error message with information to help you resolve the problem.

Logical partition attributes that change after the logical partition moves to the destination system

When you move a logical partition from one server to another, some of its attributes might change (such as the logical partition ID number) and some of its attributes remain the same (such as the logical partition configuration).

The following table describes the logical partition attributes that remain the same and the logical partition attributes that might change after you move a logical partition to the destination server.

Table 44. Logical partition attributes that might change or remain the same after a logical partition moves to the destination server

Attributes that remain the same	Attributes that might change
<ul style="list-style-type: none"> The logical partition name The logical partition type (dedicated processor or shared processor) The logical partition configuration The processor architecture The Simultaneous Multi-Threading (SMT) state of each processor The virtual MAC addresses, IP addresses, and LUN mapping to the target devices 	<ul style="list-style-type: none"> The logical partition ID number The machine type, model, and serial number The model class of the underlying server The processor version and type The processor frequency The affinity characteristics of the logical memory blocks (LMB) The maximum number of hot pluggable and installed physical processors The L1 and L2 cache size

Processor compatibility modes

Processor compatibility modes enable you to move logical partitions between servers that have different processor types without upgrading the operating environments installed in the logical partitions.

You can run several versions of the Linux and Virtual I/O Server operating environments in logical partitions on POWER5, POWER6, POWER6+, and POWER7 processor-based servers. Sometimes older versions of these operating environments do not support the capabilities that are available with new processors, thus limiting your flexibility to move logical partitions between servers that have different processor types.

A processor compatibility mode is a value assigned to a logical partition by the hypervisor that specifies the processor environment in which the logical partition can successfully operate. When you move a

logical partition to a destination server that has a different processor type from the source server, the processor compatibility mode enables that logical partition to run in a processor environment on the destination server in which it can successfully operate. In other words, the processor compatibility mode enables the destination server to provide the logical partition with a subset of processor capabilities that are supported by the operating environment that is installed in the logical partition.

Related tasks:

“Verifying the processor compatibility mode of the mobile partition” on page 66

You can use the Hardware Management Console (HMC) to determine whether the processor compatibility mode of the mobile partition is supported on the destination server, and update the mode, if necessary, so that you can successfully move the mobile partition to the destination server.

“Verifying the processor compatibility mode of the mobile partition” on page 133

You can use the Integrated Virtualization Manager (IVM) to determine whether the processor compatibility mode of the mobile partition is supported on the destination server, and update the mode, if necessary, so that you can successfully move the mobile partition to the destination server.

Processor compatibility mode definitions:

You can learn about each processor compatibility mode and the servers on which each mode can run.

The following table describes each processor compatibility mode and the servers on which the logical partitions that use each processor compatibility mode can successfully operate.

Table 45. Processor compatibility modes

Processor compatibility mode	Description	Supported servers
POWER5	The POWER5 processor compatibility mode allows you to run operating-system versions that use all the standard features of the POWER5 processor.	Logical partitions that use the POWER5 processor compatibility mode can run on POWER5, POWER6, and POWER6+ processor-based servers. Restriction: A POWER6 processor cannot emulate all features of a POWER5 processor. Similarly, a POWER7 processor cannot emulate all features of a POWER6 or a POWER5 processor. For example, certain types of performance monitoring might not be available for a logical partition if the current processor compatibility mode of a logical partition is set to the POWER5 mode.
POWER6	The POWER6 processor compatibility mode allows you to run operating-system versions that use all the standard features of the POWER6 processor.	Logical partitions that use the POWER6 processor compatibility mode can run on POWER6, POWER6+, and POWER7 processor-based servers.
POWER6+	The POWER6+ processor compatibility mode allows you to run operating-system versions that use all the standard features of the POWER6+ processor.	Logical partitions that use the POWER6+ processor compatibility mode can run on POWER6+ and POWER7 processor-based servers.

Table 45. Processor compatibility modes (continued)

Processor compatibility mode	Description	Supported servers
POWER6 enhanced	The POWER6 enhanced processor compatibility mode allows you to run operating-system versions that use all the standard features of the POWER6 processor and also provides additional floating-point instructions to applications that use the POWER6 processor.	Logical partitions that use the POWER6 enhanced processor compatibility mode can run on POWER6 processor-based servers.
POWER6+ enhanced	The POWER6+ enhanced processor compatibility mode allows you to run operating-system versions that use all the standard features of the POWER6+ processor and also provides additional floating-point instructions to applications that use the POWER6+ processor.	Logical partitions that use the POWER6+ enhanced processor compatibility mode can run on POWER6+ processor-based servers.
POWER7	The POWER7 processor compatibility mode allows you to run operating-system versions that use all the standard features of the POWER7 processor.	Logical partitions that use the POWER7 processor compatibility mode can run on POWER7 processor-based servers.
default	The default processor compatibility mode is a preferred processor compatibility mode that enables the hypervisor to determine the current mode for the logical partition. When the preferred mode is set to default, the hypervisor sets the current mode to the most fully featured mode supported by the operating environment. In most cases, this is the processor type of the server on which the logical partition is activated. For example, assume that the preferred mode is set to default and the logical partition is running on a POWER7 processor-based server. The operating environment supports the POWER7 processor capabilities so the hypervisor sets the current processor compatibility mode to POWER7.	The servers on which logical partitions with the preferred processor compatibility mode of default can run depend on the current processor compatibility mode of the logical partition. For example, if the hypervisor determines that the current mode is POWER7, the logical partition can run on POWER7 processor-based servers.

Related concepts:

“Current and preferred processor compatibility modes” on page 14

The processor compatibility mode in which the logical partition currently operates is the *current* processor compatibility mode of the logical partition. The *preferred* processor compatibility mode of a logical partition is the mode in which you want the logical partition to operate.

“Enhanced processor compatibility modes” on page 17

The POWER6 enhanced and POWER6+ enhanced processor compatibility modes provide additional floating-point instructions to applications that use the POWER6 or POWER6+ processor.

“Examples: Using processor compatibility modes in partition mobility” on page 27

You can view examples of how processor compatibility modes are used when moving an active or inactive logical partition between servers with different processor types.

Related reference:

“Migration combinations of processor compatibility modes” on page 17

View all the combinations of the processor types of the source server, the processor types of the destination server, the current and preferred processor compatibility modes of the logical partition before the migration, and the current and preferred processor compatibility modes of the logical partition after the migration.

Current and preferred processor compatibility modes:

The processor compatibility mode in which the logical partition currently operates is the *current* processor compatibility mode of the logical partition. The *preferred* processor compatibility mode of a logical partition is the mode in which you want the logical partition to operate.

The hypervisor sets the current processor compatibility mode for a logical partition by using the following information:

- The processor features supported by the operating environment running in the logical partition.
- The preferred processor compatibility mode that you specify.

When you activate the logical partition, the hypervisor checks the preferred processor compatibility mode and determines whether the operating environment supports that mode. If the operating environment supports the preferred processor compatibility mode, the hypervisor assigns the logical partition the preferred processor compatibility mode. If the operating environment does not support the preferred processor compatibility mode, the hypervisor assigns the logical partition the most fully featured processor compatibility mode that is supported by the operating environment.

The following table describes when each processor compatibility mode can be current mode or the preferred mode.

Table 46. Current and preferred processor compatibility modes

Processor compatibility mode	Can it be the current mode?	Can it be the preferred mode?
POWER5	Yes The POWER5 processor compatibility mode can be the current processor compatibility mode of a logical partition.	No You cannot specify POWER5 as the preferred processor compatibility mode. The only situation in which a logical partition will run in the POWER5 processor compatibility mode is when it is the only processor environment supported by the operating environment installed in the logical partition.
POWER6	Yes The POWER6 processor compatibility mode can be the current processor compatibility mode of a logical partition.	Yes You can specify POWER6 as the preferred processor compatibility mode for a logical partition.
POWER6+	Yes The POWER6+ processor compatibility mode can be the current processor compatibility mode of a logical partition.	Yes You can specify POWER6+ as the preferred processor compatibility mode for a logical partition.

Table 46. Current and preferred processor compatibility modes (continued)

Processor compatibility mode	Can it be the current mode?	Can it be the preferred mode?
POWER6 enhanced	Yes The POWER6 enhanced processor compatibility mode can be the current processor compatibility mode of a logical partition.	Yes You can specify POWER6 enhanced as the preferred processor compatibility mode for a logical partition.
POWER6+ enhanced	Yes The POWER6+ enhanced processor compatibility mode can be the current processor compatibility mode of a logical partition.	Yes You can specify POWER6+ enhanced as the preferred processor compatibility mode for a logical partition.
POWER7	Yes The POWER7 processor compatibility mode can be the current processor compatibility mode of a logical partition.	Yes You can specify POWER7 as the preferred processor compatibility mode for a logical partition.
default	No The default processor compatibility mode is a preferred processor compatibility mode.	Yes You can specify default as the preferred processor compatibility mode. Also, if you do not specify a preferred mode, the system automatically sets the preferred mode to default.

The following table shows the current and preferred processor compatibility modes supported on each server type.

Table 47. Processor compatibility modes supported by server type

Server processor type	Supported current modes	Supported preferred modes
POWER6+ processor-based server	POWER5, POWER6, POWER6+, POWER6+ enhanced	default, POWER6, POWER6+, POWER6+ enhanced
POWER6 processor-based server	POWER5, POWER6, POWER6 enhanced	default, POWER6, POWER6 enhanced
POWER7 processor-based server	POWER5, POWER6, POWER6+, POWER7	default, POWER6, POWER6+, POWER7

The preferred processor compatibility mode is the highest mode that the hypervisor can assign to a logical partition. If the operating environment installed in the logical partition does not support the preferred mode, the hypervisor can set the current mode to a lower mode than the preferred mode, but it cannot set the current mode to a higher mode than the preferred mode. For example, assume that a logical partition runs on a POWER7 processor-based server and you specify POWER7 as the preferred mode. The operating environment installed in the logical partition does not support the POWER7 processor capabilities, but it does support the POWER6 processor capabilities. When you activate the logical partition, the hypervisor assigns the POWER6 processor compatibility mode as the current mode for the logical partition because the POWER6 mode is the most fully featured mode that the operating environment supports and it is a lower mode than the preferred mode of POWER7.

You cannot dynamically change the current processor compatibility of a logical partition. To change the current processor compatibility mode, you must change the preferred processor compatibility mode, shut

down the logical partition, and restart the logical partition. The hypervisor attempts to set the current processor compatibility mode to the preferred mode that you specified.

When you move an active logical partition between servers with different processor types, both the current and preferred processor compatibility modes of the logical partition must be supported by the destination server. When you move an inactive logical partition between servers with different processor types, only the preferred mode of the logical partition must be supported by the destination server.

If you specify the default mode as the preferred mode for an inactive logical partition, you can move that inactive logical partition to a server of any processor type. Because all servers support the default processor compatibility mode, you can move an inactive logical partition with the preferred mode of default to a server with any processor type. When the inactive logical partition is activated on the destination server, the preferred mode remains set to default, and the hypervisor determines the current mode for the logical partition.

Related concepts:

“Examples: Using processor compatibility modes in partition mobility” on page 27

You can view examples of how processor compatibility modes are used when moving an active or inactive logical partition between servers with different processor types.

“Processor compatibility mode definitions” on page 12

You can learn about each processor compatibility mode and the servers on which each mode can run.

Related reference:

“Migration combinations of processor compatibility modes” on page 17

View all the combinations of the processor types of the source server, the processor types of the destination server, the current and preferred processor compatibility modes of the logical partition before the migration, and the current and preferred processor compatibility modes of the logical partition after the migration.

Enhanced processor compatibility modes:

The POWER6 enhanced and POWER6+ enhanced processor compatibility modes provide additional floating-point instructions to applications that use the POWER6 or POWER6+ processor.

Note: POWER7 processor-based servers do not support the enhanced mode.

If you want a logical partition to run in an enhanced mode, you must specify the enhanced mode as the preferred mode for the logical partition. If the operating environment supports the corresponding non-enhanced mode, then the hypervisor assigns the enhanced mode to the logical partition when you activate the logical partition. In other words, if you specify the POWER6+ enhanced mode as the preferred mode, and the operating environment supports the POWER6+ mode, the hypervisor assigns the POWER6+ enhanced mode to the logical partition when you activate the logical partition. Similarly, if you specify the POWER6 enhanced mode as the preferred mode, and the operating environment supports the POWER6 mode, the hypervisor assigns the POWER6 enhanced mode to the logical partition when you activate the logical partition.

Logical partitions in the POWER6 enhanced processor compatibility mode can only run on POWER6 processor-based servers, and logical partitions in the POWER6+ enhanced processor compatibility mode can only run on POWER6+ processor-based servers. Therefore, if a logical partition runs in the POWER6 enhanced mode, you can only move the logical partition to POWER6 processor-based servers. Likewise, if a logical partition runs in the POWER6+ enhanced mode, you can only move the logical partition to POWER6+ processor-based servers. If you want to move a logical partition in the POWER6 enhanced processor compatibility mode to a POWER6+ processor-based server, then you need to change the preferred mode to the default or POWER6 processor compatibility mode and restart the logical partition.

Related concepts:

“Examples: Using processor compatibility modes in partition mobility” on page 27

You can view examples of how processor compatibility modes are used when moving an active or inactive logical partition between servers with different processor types.

“Processor compatibility mode definitions” on page 12

You can learn about each processor compatibility mode and the servers on which each mode can run.

Related reference:

“Migration combinations of processor compatibility modes” on page 17

View all the combinations of the processor types of the source server, the processor types of the destination server, the current and preferred processor compatibility modes of the logical partition before the migration, and the current and preferred processor compatibility modes of the logical partition after the migration.

Migration combinations of processor compatibility modes:

View all the combinations of the processor types of the source server, the processor types of the destination server, the current and preferred processor compatibility modes of the logical partition before the migration, and the current and preferred processor compatibility modes of the logical partition after the migration.

Related concepts:

“Examples: Using processor compatibility modes in partition mobility” on page 27

You can view examples of how processor compatibility modes are used when moving an active or inactive logical partition between servers with different processor types.

“Enhanced processor compatibility modes” on page 17

The POWER6 enhanced and POWER6+ enhanced processor compatibility modes provide additional floating-point instructions to applications that use the POWER6 or POWER6+ processor.

“Current and preferred processor compatibility modes” on page 14

The processor compatibility mode in which the logical partition currently operates is the *current* processor compatibility mode of the logical partition. The *preferred* processor compatibility mode of a logical partition is the mode in which you want the logical partition to operate.

“Processor compatibility mode definitions” on page 12

You can learn about each processor compatibility mode and the servers on which each mode can run.

Migration combinations of processor compatibility modes for active partition mobility:

When you move an active logical partition between servers with different processor types, both the current and preferred processor compatibility modes of the logical partition must be supported by the destination server.

The following tables describe the processor compatibility mode combinations for active migrations. They show the processor type of the source server and the preferred and current processor compatibility modes of the logical partition on the source server before the migration. They also show the processor type of the destination server and the preferred and current processor compatibility modes of the logical partition on the destination server after the migration.

Table 48. Processor compatibility mode combinations for active migrations of POWER7 processor-based servers

Source environment			Destination environment		
Source server	Preferred mode before migration	Current mode before migration	Destination server	Preferred mode after migration	Current mode after migration
POWER7 processor-based server	default	POWER7, POWER6+, POWER6, or POWER5	POWER7 processor-based server	default	POWER7, POWER6+, POWER6, or POWER5

Table 48. Processor compatibility mode combinations for active migrations of POWER7 processor-based servers (continued)

Source environment			Destination environment		
Source server	Preferred mode before migration	Current mode before migration	Destination server	Preferred mode after migration	Current mode after migration
POWER7 processor-based server	POWER7	POWER7, POWER6+, POWER6, or POWER5	POWER7 processor-based server	POWER7	POWER7, POWER6+, POWER6, or POWER5
POWER7 processor-based server	POWER6+	POWER6+, POWER6, or POWER5	POWER7 processor-based server	POWER6+	POWER6+, POWER6, or POWER5
POWER7 processor-based server	POWER6	POWER6 or POWER5	POWER7 processor-based server	POWER6	POWER6 or POWER5
POWER7 processor-based server	default	POWER7, POWER6+, POWER6, or POWER5	POWER6+ processor-based server	default	If the current mode on the source server is POWER7, you cannot migrate the logical partition because the destination server does not support the current mode (POWER7). If the current mode on the source server is POWER6+, POWER6, or POWER5, then the current mode on the destination server is POWER6+, POWER6, or POWER5.
POWER7 processor-based server	POWER7	POWER7, POWER6+, POWER6, or POWER5	POWER6+ processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER7).	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER7).

Table 48. Processor compatibility mode combinations for active migrations of POWER7 processor-based servers (continued)

Source environment			Destination environment		
Source server	Preferred mode before migration	Current mode before migration	Destination server	Preferred mode after migration	Current mode after migration
POWER7 processor-based server	default	POWER7, POWER6+, POWER6, or POWER5	POWER6 processor-based server	default	If the current mode on the source server is POWER7 or POWER6+, you cannot migrate the logical partition because the destination server does not support the current mode (POWER7 or POWER6+). If the current mode on the source server is POWER6 or POWER5, then the current mode on the destination server is POWER6 or POWER5.
POWER7 processor-based server	POWER6+	POWER6+, POWER6, or POWER5	POWER6+ processor-based server	POWER6+	POWER6+, POWER6, or POWER5
POWER7 processor-based server	POWER6	POWER6 or POWER5	POWER6+ processor-based server	POWER6	POWER6 or POWER5
POWER7 processor-based server	POWER7 or POWER6+	POWER7, POWER6+, POWER6, or POWER5	POWER6 processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER7 or POWER6+).	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER7 or POWER6+).
POWER7 processor-based server	POWER6	POWER6 or POWER5	POWER6 processor-based server	POWER6	POWER6 or POWER5

Table 49. Processor compatibility mode combinations for active migrations of POWER6+ processor-based servers

Source environment			Destination environment		
Source server	Preferred mode before migration	Current mode before migration	Destination server	Preferred mode after migration	Current mode after migration
POWER6+ processor-based server	default	POWER6+, POWER6, or POWER5	POWER6+ processor-based server	default	POWER6+, POWER6, or POWER5

Table 49. Processor compatibility mode combinations for active migrations of POWER6+ processor-based servers (continued)

Source environment			Destination environment		
Source server	Preferred mode before migration	Current mode before migration	Destination server	Preferred mode after migration	Current mode after migration
POWER6+ processor-based server	POWER6+	POWER6+, POWER6, or POWER5	POWER6+ processor-based server	POWER6+	POWER6+, POWER6, or POWER5
POWER6+ processor-based server	POWER6+ enhanced	POWER6+ enhanced or POWER5	POWER6+ processor-based server	POWER6+ enhanced	POWER6+ enhanced or POWER5
POWER6+ processor-based server	POWER6	POWER6 or POWER5	POWER6+ processor-based server	POWER6	POWER6 or POWER5
POWER6+ processor-based server	default	POWER6+, POWER6, or POWER5	POWER6 processor-based server	default	If the current mode on the source server is POWER6+, you cannot migrate the logical partition because the destination server does not support the current mode (POWER6+). If the current mode on the source server is POWER6 or POWER5, then the current mode on the destination server is POWER6 or POWER5.
POWER6+ processor-based server	POWER6+	POWER6+, POWER6, or POWER5	POWER6 processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+).	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+).
POWER6+ processor-based server	POWER6+ enhanced	POWER6+ enhanced or POWER5	POWER6 processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+ enhanced).	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+ enhanced).

Table 49. Processor compatibility mode combinations for active migrations of POWER6+ processor-based servers (continued)

Source environment			Destination environment		
Source server	Preferred mode before migration	Current mode before migration	Destination server	Preferred mode after migration	Current mode after migration
POWER6+ processor-based server	POWER6	POWER6 or POWER5	POWER6 processor-based server	POWER6	POWER6 or POWER5
POWER6+ processor-based server	default	POWER6+, POWER6, or POWER5	POWER7 processor-based server	default	POWER7 (after you restart the logical partition), POWER6+, POWER6, or POWER5
POWER6+ processor-based server	POWER6+	POWER6+, POWER6, or POWER5	POWER7 processor-based server	POWER6+	POWER6+, POWER6, or POWER5
POWER6+ processor-based server	POWER6+ enhanced	POWER6+ enhanced or POWER5	POWER7 processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+ enhanced)	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+ enhanced)
POWER6+ processor-based server	POWER6	POWER6 or POWER5	POWER7 processor-based server	POWER6	POWER6 or POWER5

Table 50. Processor compatibility mode combinations for active migrations of POWER6 processor-based servers

Source environment			Destination environment		
Source server	Preferred mode before migration	Current mode before migration	Destination server	Preferred mode after migration	Current mode after migration
POWER6 processor-based server	default	POWER6 or POWER5	POWER6 processor-based server	default	POWER6 or POWER5
POWER6 processor-based server	POWER6	POWER6 or POWER5	POWER6 processor-based server	POWER6	POWER6 or POWER5
POWER6 processor-based server	POWER6 enhanced	POWER6 enhanced or POWER5	POWER6 processor-based server	POWER6 enhanced	POWER6 enhanced or POWER5
POWER6 processor-based server	default	POWER6 or POWER5	POWER6+ processor-based server	default	POWER6+ (after you restart the logical partition), POWER6, or POWER5
POWER6 processor-based server	POWER6	POWER6 or POWER5	POWER6+ processor-based server	POWER6	POWER6 or POWER5

Table 50. Processor compatibility mode combinations for active migrations of POWER6 processor-based servers (continued)

Source environment			Destination environment		
Source server	Preferred mode before migration	Current mode before migration	Destination server	Preferred mode after migration	Current mode after migration
POWER6 processor-based server	POWER6 enhanced	POWER6 enhanced or POWER5	POWER6+ processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6 enhanced).	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6 enhanced).
POWER6 processor-based server	default	POWER6 or POWER5	POWER7 processor-based server	default	POWER7 (after you restart the logical partition), POWER6 or POWER5
POWER6 processor-based server	POWER6	POWER6 or POWER5	POWER7 processor-based server	POWER6	POWER6 or POWER5
POWER6 processor-based server	POWER6 enhanced	POWER6 enhanced or POWER5	POWER7 processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6 enhanced)	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6 enhanced)

Related reference:

“Migration combinations of processor compatibility modes for inactive partition mobility” on page 23
 When you move an inactive logical partition between servers with different processor types, only the preferred mode of the logical partition must be supported by the destination server.

“Migration combinations for version 1.5, and earlier, of the IVM” on page 114
 Learn about the processor compatibility mode combinations for migrations where versions 1.5 (and earlier) of the Integrated Virtualization Manager (IVM) manage the source server and versions 2.1 (and later) of the IVM manage the destination server.

Migration combinations of processor compatibility modes for inactive partition mobility:

When you move an inactive logical partition between servers with different processor types, only the preferred mode of the logical partition must be supported by the destination server.

The following tables describe the processor compatibility mode combinations for inactive migrations. They show the processor type of the source server and the preferred processor compatibility modes of the logical partition on the source server before the migration. They also show the processor type of the destination server and the preferred and current processor compatibility modes of the logical partition on the destination server after the migration.

Table 51. Processor compatibility mode combinations for inactive migrations of POWER7 processor-based servers

Source environment		Destination environment		
Source server	Preferred mode before migration	Destination server	Preferred mode before migration	Current mode after migration
POWER7 processor-based server	default	POWER7 processor-based server	default	POWER7, POWER6+, POWER6, or POWER5
POWER7 processor-based server	POWER7	POWER7 processor-based server	POWER7	POWER7, POWER6+, POWER6, or POWER5
POWER7 processor-based server	POWER6+	POWER7 processor-based server	POWER6+	POWER6+, POWER6, or POWER5
POWER7 processor-based server	POWER6	POWER7 processor-based server	POWER6	POWER6 or POWER5
POWER7 processor-based server	default	POWER6+ processor-based server	default	POWER6+, POWER6, or POWER5
POWER7 processor-based server	POWER6+	POWER6+ processor-based server	POWER6+	POWER6+, POWER6, or POWER5
POWER7 processor-based server	POWER6	POWER6+ processor-based server	POWER6	POWER6 or POWER5
POWER7 processor-based server	POWER7	POWER6+ processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER7).	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER7).
POWER7 processor-based server	default	POWER6 processor-based server	default	POWER6 or POWER5
POWER7 processor-based server	POWER7 or POWER6+	POWER6 processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER7 or POWER6+).	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER7 or POWER6+).
POWER7 processor-based server	POWER6	POWER6 processor-based server	POWER6	POWER6 or POWER5

Table 52. Processor compatibility mode combinations for inactive migrations of POWER6+ processor-based servers

Source environment		Destination environment		
Source server	Preferred mode before migration	Destination server	Preferred mode before migration	Current mode after migration
POWER6+ processor-based server	default	POWER6+ processor-based server	default	POWER6+, POWER6, or POWER5
POWER6+ processor-based server	POWER6+	POWER6+ processor-based server	POWER6+	POWER6+, POWER6, or POWER5
POWER6+ processor-based server	POWER6	POWER6+ processor-based server	POWER6	POWER6 or POWER5
POWER6+ processor-based server	POWER6+ enhanced	POWER6+ processor-based server	POWER6+ enhanced	POWER6+ enhanced or POWER5
POWER6+ processor-based server	default	POWER6 processor-based server	default	POWER6 or POWER5
POWER6+ processor-based server	POWER6+	POWER6 processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+).	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+).
POWER6+ processor-based server	POWER6	POWER6 processor-based server	POWER6	POWER6 or POWER5
POWER6+ processor-based server	POWER6+ enhanced	POWER6 processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+ enhanced).	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+ enhanced).
POWER6+ processor-based server	default	POWER7 processor-based server	default	POWER7 (after you restart the logical partition), POWER6+, POWER6, or POWER5
POWER6+ processor-based server	POWER6+	POWER7 processor-based server	POWER6+	POWER6+, POWER6, or POWER5
POWER6+ processor-based server	POWER6+ enhanced	POWER7 processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+ enhanced)	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6+ enhanced)

Table 52. Processor compatibility mode combinations for inactive migrations of POWER6+ processor-based servers (continued)

Source environment		Destination environment		
Source server	Preferred mode before migration	Destination server	Preferred mode before migration	Current mode after migration
POWER6+ processor-based server	POWER6	POWER7 processor-based server	POWER6	POWER6 or POWER5

Table 53. Processor compatibility mode combinations for inactive migrations of POWER6 processor-based servers

Source environment		Destination environment		
Source server	Preferred mode before migration	Destination server	Preferred mode before migration	Current mode after migration
POWER6 processor-based server	default	POWER6 processor-based server	default	POWER6 or POWER5
POWER6 processor-based server	POWER6	POWER6 processor-based server	POWER6	POWER6 or POWER5
POWER6 processor-based server	POWER6 enhanced	POWER6 processor-based server	POWER6 enhanced	POWER6 enhanced or POWER5
POWER6 processor-based server	default	POWER6+ processor-based server	default	POWER6+, POWER6, or POWER5
POWER6 processor-based server	POWER6	POWER6+ processor-based server	POWER6	POWER6 or POWER5
POWER6 processor-based server	POWER6 enhanced	POWER6+ processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6 enhanced).	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6 enhanced).
POWER6 processor-based server	default	POWER7 processor-based server	default	POWER7 (after you restart the logical partition), POWER6, or POWER5
POWER6 processor-based server	POWER6	POWER7 processor-based server	POWER6	POWER6 or POWER5
POWER6 processor-based server	POWER6 enhanced	POWER7 processor-based server	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6 enhanced)	You cannot migrate the logical partition because the destination server does not support the preferred mode (POWER6 enhanced)

Related reference:

“Migration combinations of processor compatibility modes for active partition mobility” on page 18
 When you move an active logical partition between servers with different processor types, both the current and preferred processor compatibility modes of the logical partition must be supported by the destination server.

“Migration combinations for version 1.5, and earlier, of the IVM”

Learn about the processor compatibility mode combinations for migrations where versions 1.5 (and earlier) of the Integrated Virtualization Manager (IVM) manage the source server and versions 2.1 (and later) of the IVM manage the destination server.

Migration combinations for version 1.5, and earlier, of the IVM:

Learn about the processor compatibility mode combinations for migrations where versions 1.5 (and earlier) of the Integrated Virtualization Manager (IVM) manage the source server and versions 2.1 (and later) of the IVM manage the destination server.

The following table shows the processor type of the source server and the processor compatibility mode of the logical partition on the source server before the migration. It also shows the processor type of the destination server and the preferred and current processor compatibility modes of the logical partition on the destination server after the migration.

Table 54. Processor compatibility mode combinations for mixed versions of the IVM

Source environment		Destination environment		
Source server	Mode before migration	Destination server	Preferred mode after migration	Current mode after migration
POWER6 processor-based server	default	POWER6 processor-based server	POWER6	POWER6
POWER6 processor-based server	POWER6 enhanced	POWER6 processor-based server	POWER6 enhanced	POWER6 enhanced or POWER6
POWER6 processor-based server	default	POWER6+ processor-based server	POWER6	POWER6
POWER6 processor-based server	POWER6 enhanced	POWER6+ processor-based server	You cannot migrate the logical partition because the destination server does not support the enhanced mode.	You cannot migrate the logical partition because the destination server does not support the enhanced mode.

Requirement: The previous table does not list POWER6+ processor-based servers or POWER7 processor-based servers as the source server. If you plan to manage a POWER6+ processor-based server with the IVM, the IVM must be at version 2.1, or later. If you plan to manage a POWER7 processor-based server with the IVM, the IVM must be at version 2.1.2 with fix pack 22.1 and service pack 1, or later. If you plan to move a logical partition from a POWER6 or POWER6+ processor-based server to a POWER7 processor-based server, the IVM that manages the POWER6 or POWER6+ processor-based server must be at version 2.1.2 with fix pack 22, or later.

Related reference:

“Migration combinations of processor compatibility modes for active partition mobility” on page 18
 When you move an active logical partition between servers with different processor types, both the current and preferred processor compatibility modes of the logical partition must be supported by the destination server.

“Migration combinations of processor compatibility modes for inactive partition mobility” on page 23
 When you move an inactive logical partition between servers with different processor types, only the

preferred mode of the logical partition must be supported by the destination server.

Examples: Using processor compatibility modes in partition mobility:

You can view examples of how processor compatibility modes are used when moving an active or inactive logical partition between servers with different processor types.

Moving an active logical partition from a POWER6 processor-based server to a POWER7 processor-based server

You want to move an active logical partition from a POWER6 processor-based server to a POWER7 processor-based server so that the logical partition can take advantage of the additional capabilities available with the POWER7 processor.

To accomplish this task, perform the following steps:

1. Set the preferred processor compatibility mode to the default mode. When you activate the logical partition on the POWER6 processor-based server, it runs in the POWER6 mode.
2. Move the logical partition to the POWER7 processor-based server. Both the current and preferred modes remain unchanged for the logical partition until you restart the logical partition.
3. Restart the logical partition on the POWER7 processor-based server. The hypervisor evaluates the configuration. Because the preferred mode is set to default and the logical partition now runs on a POWER7 processor-based server, the highest mode available is the POWER7 mode. The hypervisor determines that the most fully featured mode supported by the operating environment installed in the logical partition is the POWER7 mode and changes the current mode of the logical partition to the POWER7 mode.

At this point, the current processor compatibility mode of the logical partition is the POWER7 mode and the logical partition runs on the POWER7 processor-based server.

Moving the active logical partition back to the POWER6 processor-based server

A problem arises and you need to move the active logical partition back to the POWER6 processor-based server. Because the logical partition now runs in the POWER7 mode and the POWER7 mode is not supported on the POWER6 processor-based server, you need to adjust the preferred mode for the logical partition so that the hypervisor can reset the current mode to a mode that is supported by the POWER6 processor-based server.

To move the logical partition back to the POWER6 processor-based server, perform the following steps:

1. Change the preferred mode from the default mode to the POWER6 mode.
2. Restart the logical partition on the POWER7 processor-based server. The hypervisor evaluates the configuration. Because the preferred mode is set to POWER6, the hypervisor does not set the current mode to a higher mode than POWER6. Remember, the hypervisor first determines whether it can set the current mode to the preferred mode. If not, it determines whether it can set the current mode to the next highest mode, and so on. In this case, the operating environment supports the POWER6 mode, so the hypervisor sets the current mode to the POWER6 mode.
3. Now that the logical partition runs in the POWER6 mode and the POWER6 mode is supported on the POWER6 processor-based server, move the logical partition back to the POWER6 processor-based server.

Moving an active logical partition between different processor types without making configuration changes

Depending on how often you need to move logical partitions, you might want to maintain the flexibility to move an active logical partition between a POWER6 processor-based server and a POWER7 processor-based server so that you can move the logical partition back and forth without making

configuration changes. The easiest way to maintain this type of flexibility is to determine the processor compatibility mode supported on both the source and destination servers and set the preferred processor compatibility mode of the logical partition to the highest mode supported by both servers.

To accomplish this flexibility, perform the following steps:

1. Set the preferred processor compatibility mode to the POWER6 mode because the POWER6 mode is the highest mode supported by both POWER6 processor-based servers and POWER7 processor-based servers.
2. Move the logical partition from the POWER6 processor-based server to the POWER7 processor-based server.
3. Restart the logical partition on the POWER7 processor-based server. The hypervisor evaluates the configuration. Remember, the hypervisor does not set the current mode higher than the preferred mode. First, the hypervisor determines whether it can set the current mode to the preferred mode. If not, it then determines whether it can set the current mode to the next highest mode, and so on. In this case, the operating environment supports the POWER6 mode, so the hypervisor sets the current mode to the POWER6 mode.
4. Do not make any configuration changes to move the logical partition back to the POWER6 processor-based server because the POWER6 mode is supported on the POWER6 processor-based server.
5. Move the logical partition back to the POWER6 processor-based server.
6. Restart the logical partition on the POWER6 processor-based server. The hypervisor evaluates the configuration. The hypervisor determines that the operating environment supports the preferred mode of POWER6, and sets the current mode to the POWER6 mode.

Moving an inactive logical partition between servers with different processor types

The same logic from the previous examples applies to inactive partition mobility, except inactive partition mobility does not need the current processor compatibility mode of the logical partition because the logical partition is inactive. After you move an inactive logical partition to the destination server and activate that logical partition on the destination server, the hypervisor evaluates the configuration and sets the current mode for the logical partition just like it does when you restart a logical partition after active partition mobility. The hypervisor attempts to set the current mode to the preferred mode. If it cannot, it checks the next highest mode, and so on.

Related concepts:

“Enhanced processor compatibility modes” on page 17

The POWER6 enhanced and POWER6+ enhanced processor compatibility modes provide additional floating-point instructions to applications that use the POWER6 or POWER6+ processor.

“Current and preferred processor compatibility modes” on page 14

The processor compatibility mode in which the logical partition currently operates is the *current* processor compatibility mode of the logical partition. The *preferred* processor compatibility mode of a logical partition is the mode in which you want the logical partition to operate.

“Processor compatibility mode definitions” on page 12

You can learn about each processor compatibility mode and the servers on which each mode can run.

Related reference:

“Migration combinations of processor compatibility modes” on page 17

View all the combinations of the processor types of the source server, the processor types of the destination server, the current and preferred processor compatibility modes of the logical partition before the migration, and the current and preferred processor compatibility modes of the logical partition after the migration.

Partition mobility environment

You can learn about each component of the partition mobility environment and its contribution in enabling successful partition mobility. Components of the partition mobility environment include the

source and destination servers, the Integrated Virtualization Manager (IVM), the mobile partition, the networking configuration, and the storage configuration.

Related tasks:

“Preparing for partition mobility” on page 123

You need to verify that the source and destination systems are configured correctly so that you can successfully move the mobile partition from the source system to the destination system. This includes verifying the configuration of the source and destination servers, the Integrated Virtualization Manager (IVM) management partitions, the mobile partition, the virtual storage configuration, and the virtual network configuration.

Source and destination servers in a partition mobility environment:

Two servers are involved in partition mobility that is managed by the Integrated Virtualization Manager (IVM). The *source server* is the server from which you want to move the logical partition, and the *destination server* is the server to which you want to move the logical partition.

The source and destination servers must be POWER6 processor-based servers, or later, to participate in partition mobility. The destination server must have enough available processor and memory resources to allow the mobile partition to run on its server.

Shared memory is physical memory that is assigned to the shared memory pool and shared among multiple logical partitions. The *shared memory pool* is a defined collection of physical memory blocks that are managed as a single memory pool by the hypervisor. Logical partitions that you assign to the shared memory pool share the memory in the pool with other logical partitions that you assign to the pool.


If the mobile partition uses shared memory on the source server, the destination server must also have a shared memory pool to which the mobile partition can be assigned. If the mobile partition uses dedicated memory on the source server, it must also use dedicated memory on the destination server.

Related tasks:

“IVM-managed systems: Preparing the source and destination servers for partition mobility” on page 123

You need to verify that the source and destination servers are configured correctly so that you can successfully move the mobile partition from the source server to the destination server by using the Integrated Virtualization Manager (IVM). This includes tasks such as verifying the logical memory block size of the source and destination servers, and verifying the available memory and processor resources of the destination server.

Related information:

 Overview of shared memory

 Stopping a Dynamic Platform Optimizer operation

Integrated Virtualization Manager in a partition mobility environment:

Learn about the Integrated Virtualization Manager (IVM) and how you can use it to move an active or inactive logical partition from one server to another server.

When you install the Virtual I/O Server on a system that is not managed by an HMC or an IBM BladeCenter blade server, the Virtual I/O Server becomes the management partition and provides the IVM for systems management. The IVM provides a web-based and command-line interface that you can use to migrate a logical partition from one system to another.

The Migration task on the IVM helps you validate and complete a partition migration. The IVM determines the appropriate type of migration to use based on the state of the logical partition. If the logical partition is in the *Running* state, then the migration is active. If the logical partition is in the *Not Activated* state, then the migration is inactive. Before migrating your logical partition, conduct a validation check to ensure that your migration will complete successfully.

The following table describes the services that the management partitions on the source and destination servers provide to the mobile partition (and other client partitions).

Table 55. Services provided by the management partition

Service provided by the management partitions	Description
Server partition	<p>The management partition on the source server and the management partition on the destination server must provide storage and networking resources to the mobile partition so that the mobile partition has access to the same storage from both the source and destination servers.</p> <p>Where possible, partition mobility preserves the following configuration attributes:</p> <ul style="list-style-type: none"> • User-defined names of virtual target devices. • User-defined adapter IDs of the virtual server adapters.
Mover service partition	<p>For active partition mobility, the management partition on the source server and the management partition on the destination server automatically become mover service partitions. During active partition mobility, the mover service partitions transfer the mobile partition from the source server to the destination server as follows:</p> <ul style="list-style-type: none"> • On the source server, the mover service partition extracts the logical partition state information of the mobile partition from the hypervisor. • The mover service partition on the source server sends the logical partition state information to the mover service partition on the destination server. • On the destination server, the mover service partition installs the logical partition state information to the hypervisor.
Paging VIOS partition	<p>A VIOS logical partition that is assigned to the shared memory pool (hereafter referred to as a <i>paging VIOS partition</i>) provides access to the paging space devices for the logical partitions that use shared memory. The management partition on the source server is the paging VIOS partition on the source server, and the management partition on the destination server is the paging VIOS partition on the destination server.</p> <p>When you validate a mobile partition (that uses shared memory) for active partition mobility, the IVM checks that the paging storage pool on the destination system contains an available paging space device that meets the size requirements of the mobile partition. If the paging storage pool does not contain such a device, the IVM checks that the paging storage pool has enough space for it to automatically create a paging space device that meets the size requirements of the mobile partition.</p>

Related concepts:

“Network configuration in a partition mobility environment” on page 119

In partition mobility that is managed by the Integrated Virtualization Manager (IVM), the network between the source and destination servers is used to pass the mobile partition state information and other configuration data from the source environment to the destination environment. The mobile

partition uses the virtual LAN for network access.

“Preparing the source and destination management partitions for partition mobility” on page 129
You need to verify that the source and destination management partitions are configured correctly so that you can successfully move the mobile partition from the source server to the destination server. This includes verifying the version of the Integrated Virtualization Manager (IVM) and activating the PowerVM hardware feature for IBM PowerLinux servers.

Software applications that recognize partition mobility:

Software applications might be designed to recognize and adapt to changes in the system hardware after being moved from one system to another.

Most software applications running in Linux logical partitions will not require any changes to work correctly during active partition mobility. Some applications might have dependencies on characteristics that change between the source and destination servers and other applications might need to adjust to support the migration.

PowerHA (or High Availability Cluster Multi-Processing) clustering software is aware of partition mobility. You can move a mobile partition that is running the PowerHA clustering software to another server without restarting the PowerHA clustering software.

Examples of applications that would benefit if they were aware of partition mobility:

- Software applications that use processor and memory affinity characteristics to tune their behavior because affinity characteristics might change as a result of migration. The application's functions remain the same, but performance variations may be observed.
- Applications that use processor binding will maintain their binding to the same logical processors across migrations, but in reality the physical processors will change. Binding is usually done to maintain hot caches, but the physical processor move operation will require a cache hierarchy on the destination system. This usually occurs very quickly and should not be visible to the users.
- Applications that are tuned for given cache architectures, such as hierarchy, size, line-size, and associativity. These applications are usually limited to high-performance computing applications, but the just-in-time (JIT) compiler of the Java™ Virtual Machine is also optimized for the cache-line size of the processor on which it was opened.
- Performance analysis, capacity planning, and accounting tools and their agents are usually migration-aware because the processor performance counters might change between the source and destination servers, as might the processor type and frequency. Additionally, tools that calculate an aggregate system load based on the sum of the loads in all hosted logical partitions must be aware that a logical partition has left the system or that a new logical partition arrived.
- Workload managers

Network configuration in a partition mobility environment:

In partition mobility that is managed by the Integrated Virtualization Manager (IVM), the network between the source and destination servers is used to pass the mobile partition state information and other configuration data from the source environment to the destination environment. The mobile partition uses the virtual LAN for network access.

During active partition mobility, it is important that the two management partitions must be able to communicate with each other. The virtual LAN must be bridged to a physical network using a virtual Ethernet bridge in the management partition. The LAN must be configured so that the mobile partition can continue to communicate with other necessary clients and servers after a migration is completed.

Active partition mobility has no specific requirements on the memory size of the mobile partition. The memory transfer is a procedure that does not interrupt the activity of the mobile partition and might take

time when a large memory configuration is involved on a slow network. Because of this reason, use a high-bandwidth connection, such as Gigabit Ethernet. The network bandwidth between the mover service partitions must be 1 GB or greater.

With VIOS 2.1.2.0, or later, you can enable secure IP tunnels between the mover service partition on the source server and the mover service partition on the destination server. For example, you might want to enable secure IP tunnels when the source and destination servers are not on a trusted network. Secure IP tunnels encrypt the partition state information that the mover service partitions exchange during active partition mobility. Mover service partitions with secure IP tunnels might require slightly more processing resources.

Related concepts:

“Integrated Virtualization Manager in a partition mobility environment” on page 117

Learn about the Integrated Virtualization Manager (IVM) and how you can use it to move an active or inactive logical partition from one server to another server.

Related tasks:

“Preparing the network configuration for partition mobility” on page 134

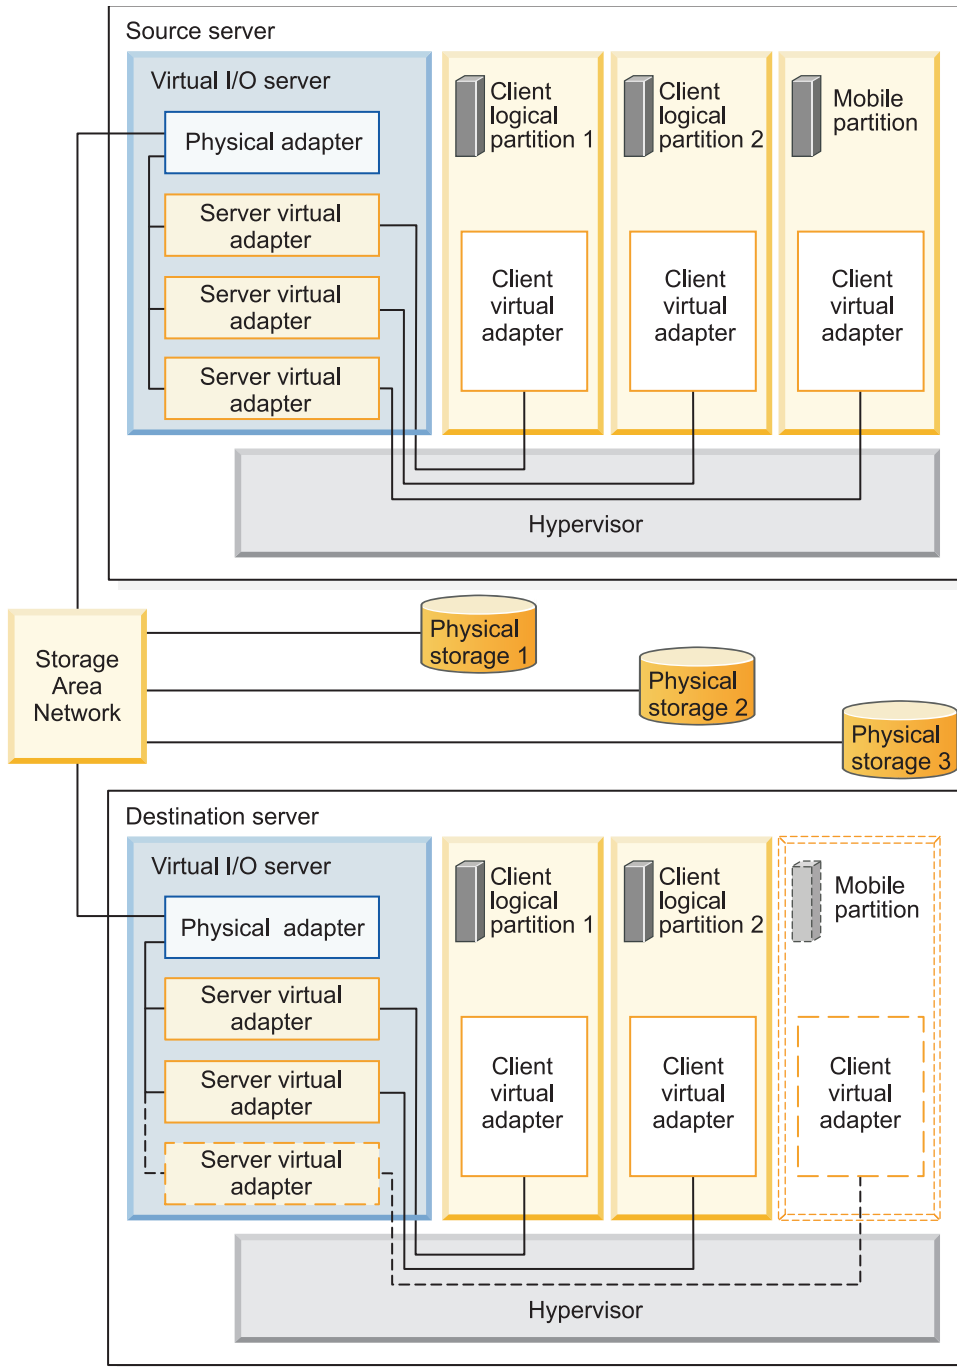
You need to verify that the network configuration is configured correctly so that you can successfully move the mobile partition from the source server to the destination server by using the Integrated Virtualization Manager (IVM). This includes tasks such as configuring a virtual Ethernet bridge on the source and destination management partitions and creating at least one virtual Ethernet adapter on the mobile partition.

Storage configuration in a partition mobility environment:

Learn about the virtual SCSI and virtual Fibre Channel configuration required for partition mobility that is managed by the Integrated Virtualization Manager (IVM).

The mobile partition moves from one server to another by the source server sending the logical partition state information to the destination server over a local area network (LAN). However, partition disk data cannot pass from one system to another system over a network. Thus, for partition mobility to succeed, the mobile partition must use storage resources that are managed by a storage area network (SAN). By using SAN storage, the mobile partition can access the same storage from both the source and destination servers.

The following figure shows an example of the storage configuration required for partition mobility.



The physical storage that the mobile partition uses, Physical storage 3, is connected to the SAN. At least one physical adapter that is assigned to the source Virtual I/O Server management partition is connected to the SAN, and at least one physical adapter that is assigned to the destination Virtual I/O Server management partition is also connected to the SAN.

If the mobile partition connects to Physical storage 3 through virtual Fibre Channel adapters, the physical adapters that are assigned to the source and destination Virtual I/O Server management partitions must support N_Port ID Virtualization (NPIV).

The physical adapter on the source Virtual I/O Server management partition connects to one or more virtual adapters on the source Virtual I/O Server management partition. Similarly, the physical adapter on the destination Virtual I/O Server management partition connects to one or more virtual adapters on

the destination Virtual I/O Server management partition. If the mobile partition connects to Physical storage 3 through virtual SCSI adapters, the virtual adapters on both the source and destination Virtual I/O Server management partitions are assigned to access the logical unit numbers (LUNs) of Physical storage 3.

Each virtual adapter on the source Virtual I/O Server management partition connects to at least one virtual adapter on a client logical partition. Similarly, each virtual adapter on the destination Virtual I/O Server management partition connects to at least one virtual adapter on a client logical partition.

Each virtual Fibre Channel adapter that is created on the mobile partition (or any client logical partition) is assigned a pair of worldwide port names (WWPNs). Both WWPNs in the WWPN pair are assigned to access the LUNs of the physical storage that the mobile partition uses, or Physical storage 3. During normal operation, the mobile partition uses one WWPN to log on to the SAN and access Physical Storage 3. When you move the mobile partition to the destination server, there is a brief period where the mobile partition runs on both the source and destination servers. Because the mobile partition cannot log on to the SAN from both the source and destination servers at the same time using the same WWPN, the mobile partition uses the second WWPN to log on to the SAN from the destination server during the migration. The WWPNs of each virtual Fibre Channel adapter move with the mobile partition to the destination server.

When you move the mobile partition to the destination server, the IVM (that manages the destination server) performs the following tasks on the destination server:

- Creates virtual adapters on the destination Virtual I/O Server logical partition
- Connects the virtual adapters on the destination Virtual I/O Server logical partition to the virtual adapters on the mobile partition

Important: The IVM automatically creates and manages the virtual adapters previously described. The IVM automatically adds and removes virtual SCSI adapters to and from the management partition and the logical partitions when you create and delete a logical partition. The IVM automatically adds and removes virtual Fibre Channel adapters to and from the management partition and the logical partitions when you assign and unassign logical partitions to and from physical Fibre Channel ports using the graphical user interface.

Related concepts:

“Integrated Virtualization Manager in a partition mobility environment” on page 117

Learn about the Integrated Virtualization Manager (IVM) and how you can use it to move an active or inactive logical partition from one server to another server.

Related tasks:

“Preparing the virtual SCSI configuration for partition mobility” on page 137

You need to verify that the virtual SCSI configuration is configured correctly so that you can successfully move the mobile partition from the source server to the destination server by using the Integrated Virtualization Manager (IVM). This includes tasks such as verifying the `reserve_policy` of the physical volumes, and verifying that the virtual devices have the same unique identifier, physical identifier, or IEEE volume attribute.

“Preparing the virtual Fiber Channel configuration for partition mobility” on page 140

You need to verify that the virtual Fiber Channel configuration is configured correctly so that you can successfully move the mobile partition from the source server to the destination server by using the Integrated Virtualization Manager (IVM). The verification includes tasks such as verifying the worldwide port names (WWPNs) of the virtual Fiber Channel adapters on the mobile partition, and verifying that the physical Fibre Channel adapters and the physical Fiber Channel switches support NPIV.

Related information:

 [Redundancy configuration using virtual fibre channel adapters](#)

Preparing for partition mobility

You need to verify that the source and destination systems are configured correctly so that you can successfully move the mobile partition from the source system to the destination system. This includes verifying the configuration of the source and destination servers, the Integrated Virtualization Manager (IVM) management partitions, the mobile partition, the virtual storage configuration, and the virtual network configuration.

Related concepts:

“Partition mobility overview for IVM” on page 94

You can learn about the benefits of partition mobility, how the Integrated Virtualization Manager (IVM) performs active and inactive partition mobility, and about the configuration that is required to successfully move a logical partition from one system to another.

“Partition mobility environment” on page 116

You can learn about each component of the partition mobility environment and its contribution in enabling successful partition mobility. Components of the partition mobility environment include the source and destination servers, the Integrated Virtualization Manager (IVM), the mobile partition, the networking configuration, and the storage configuration.

IVM-managed systems: Preparing the source and destination servers for partition mobility

You need to verify that the source and destination servers are configured correctly so that you can successfully move the mobile partition from the source server to the destination server by using the Integrated Virtualization Manager (IVM). This includes tasks such as verifying the logical memory block size of the source and destination servers, and verifying the available memory and processor resources of the destination server.

To prepare the source and destination servers for active or inactive partition mobility, complete the following tasks.

Table 56. Preparation tasks for the source and destination servers

Server planning tasks	Active mobility task	Inactive mobility task	Information resources
<p>1. Ensure that the source and destination servers are one of the following POWER7 models:</p> <ul style="list-style-type: none"> • 8246-L1C • 8246-L1D • 8246-L1S • 8246-L1T • 8246-L2C • 8246-L2D • 8246-L2S • 8246-L2T <p>Notes:</p> <ul style="list-style-type: none"> • The source and destination servers can also be POWER6 processor-based servers. See “Processor compatibility mode definitions” on page 12 for processor compatibility mode information. • Ensure that the destination server has the necessary software licenses and support maintenance contracts. To verify the entitlements that are active on your servers, see the Entitled Software Support website. 	X	X	
2. Ensure that the firmware levels on the source and destination servers are compatible.	X	X	“Firmware support matrix for partition mobility” on page 51

Table 56. Preparation tasks for the source and destination servers (continued)

Server planning tasks	Active mobility task	Inactive mobility task	Information resources
3. Ensure that the logical memory block size is the same on the source and destination servers. Determine the logical memory block size of each server, and update the sizes if necessary.	X	X	Viewing and modifying system properties
4. If the mobile partition uses shared memory, ensure that the shared memory pool is created on the destination server.	X	X	Defining the shared memory pool by using the Integrated Virtualization Manager
5. Ensure that the destination server has enough available memory to support the mobile partition.	X	X	<ul style="list-style-type: none"> • If the mobile partition uses dedicated memory, see “Determining the available physical memory on the destination server” on page 126. • If the mobile partition uses shared memory, see “Determining the available I/O entitled memory on the destination server” on page 126.
6. Ensure that the destination server has enough available processors to support the mobile partition.	X	X	“Determining available processors on the destination server” on page 128
7. Verify that the source and destination Virtual I/O Server can communicate with each other.	X	X	
<p>8. If the mobile partition on the source server is suspend-capable, verify that the destination server also supports suspend-capable partitions. You must also verify that there is a minimum of one reserved storage device with a size of at least 110% of the maximum partition memory.</p> <p>Before you can use a disk from the reserved storage pool, you might need to clear the first 4096 bytes of the disk when you see a message that indicates that the disk cannot be used. The disk might have stale data indicating that it is already in use by another partition on the managed system, or might actively be used by another managed system. You must verify with your system administrator whether the disk is currently being used. You must initialize the first 4096 bytes of the disk only after you have ensured that the disk is no longer in use and that you have corrected all of the configuration issues associated with using the disk.</p>	X	X	

Related concepts:

“Source and destination servers in a partition mobility environment” on page 117

Two servers are involved in partition mobility that is managed by the Integrated Virtualization Manager (IVM). The *source server* is the server from which you want to move the logical partition, and the *destination server* is the server to which you want to move the logical partition.

Partition mobility firmware support matrix:

Ensure that the firmware levels on the source and destination servers are compatible before upgrading.

In the following table, the values in the left column represent the firmware level you are migrating from, and the values in the top row represent the firmware level you are migrating to. For each combination, *blocked* entries are blocked by code from migrating. *Not supported* entries are not blocked from migrating, but are not supported by IBM. *Mobile* entries are eligible for migration.

Table 57. Firmware level

Migrating from firmware level	350_xxx, or later	710_xxx	720_xxx	730_xxx	740_xxx	760_xxx	763_xxx	770_xxx	773_xxx	780_xxx
340_039, or later	Mobile	Mobile	Mobile	Blocked	Blocked	Blocked	Blocked	Blocked	Blocked	Blocked
350_xxx, or later	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile
710_xxx	Mobile	Mobile	Mobile	Mobile	Mobile	Blocked	Blocked	Blocked	Blocked	Blocked
720_xxx	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Blocked	Blocked	Blocked
730_xxx	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile
740_xxx	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile
760_xxx	Mobile	Blocked	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile
763_xxx	Mobile	Blocked	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile
770_xxx	Mobile	Blocked	Blocked	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile
773_xxx	Mobile	Blocked	Blocked	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile
780_xxx	Mobile	Blocked	Blocked	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile	Mobile

The following table shows the number of concurrent migrations that are supported per system. The corresponding minimum levels of firmware and Virtual I/O Server (VIOS) that are required are also shown.

Table 58. Concurrent migrations

Concurrent migrations per system	Firmware level	VIOS level	Maximum concurrent migrations per VIOS
4	All	All	4
10	Level 7.6, or later	Version 2.2.2.0	8

Restrictions:

- All concurrent migrations must have the same source and target system.
- Systems that are managed by the Integrated Virtualization Manager (IVM) support up to ten concurrent migrations.
- You can run up to four concurrent Suspend/Resume operations.
- You cannot perform Live Partition Mobility that is both bidirectional and concurrent. For example:
 - When you are moving a mobile partition from the source server to the destination server, you cannot move another mobile partition from the destination server to the source server.
 - When you are moving a mobile partition from the source server to the destination server, you cannot move another mobile partition from the destination server to some other server.

Determining the available physical memory on the destination server:

You can determine whether the destination server has enough physical memory available to support the mobile partition. You can then make more physical memory available, if necessary, by using the Integrated Virtualization Manager (IVM).

Use any role other than View Only to perform this task. Users with the Service Representative (SR) user role cannot view or modify storage values.

To determine whether the destination server has enough physical memory available to support the mobile partition, complete the following steps from the IVM:

1. Identify the amount of physical memory that the mobile partition requires:
 - a. From the **Partition Management** menu, click **View/Modify Partitions**. The View/Modify Partition panel is displayed.
 - b. Select the mobile partition.
 - c. From the **Tasks** menu, click **Properties**. The Partition Properties window is displayed.
 - d. Click the **Memory** tab.
 - e. Record the minimum, assigned, and maximum memory settings.
 - f. Click **OK**
2. Identify the amount of physical memory that is available on the destination server:
 - a. From the **Partition Management** menu, click **View/Modify System Properties**. The View/Modify System Properties window is displayed.
 - b. Click the **Memory** tab.
 - c. From the **General** section, record the **Current memory available** and the **Reserved firmware memory**.
3. Compare the values from steps 1 and 2.

Keep in mind that when you move the mobile partition to the destination server, the destination server requires more reserved firmware memory to manage the mobile partition. If the destination server does not have enough physical memory available to support the mobile partition, you can add more available physical memory to the destination server by performing one or more of the following tasks:

 - Dynamically remove physical memory from logical partitions that use dedicated memory. For instructions, see *Dynamically managing memory*.
 - If the destination server is configured with a shared memory pool, dynamically remove physical memory from the shared memory pool. For instructions, see *Changing the shared memory pool size by using the Integrated Virtualization Manager*.

Determining the available I/O entitled memory on the destination server:

You can determine whether the shared memory pool on the destination server has enough available memory to accommodate the I/O entitled memory required by the mobile partition. You can then allocate more physical memory to the shared memory pool, if necessary, by using the Integrated Virtualization Manager (IVM).

To determine whether the shared memory pool on the destination server has enough available memory to accommodate the I/O entitled memory required by the mobile partition, complete the following steps from the IVM:

1. Identify the amount of I/O entitled memory that the mobile partition requires:
 - a. In the navigation pane, click **View/Modify Partitions** under **Partition Management**. The View/Modify Partitions page is displayed.
 - b. Select the mobile partition.

- c. From the **Tasks** menu, click **Properties**. The Partition Properties page is displayed.
 - d. Click the **Memory** tab.
 - e. Record the **I/O entitled memory**.
2. Identify the amount of available physical memory in the shared memory pool on the destination server:
 - a. In the navigation pane, click **View/Modify Shared Memory Pool** under **Partition Management**. The View/Modify System Properties page is displayed.
 - b. Note the amount of available memory shown in the **Shared memory pool size** field.
 3. Compare the amount of available memory (from step 2) with the amount of I/O entitled memory required by the mobile partition (from step 1).
 - If more memory is available than the amount of I/O entitled memory required by the mobile partition, the shared memory pool on the destination server has enough available memory to support the mobile partition on the destination server.
 - If the amount of I/O entitled memory required by the mobile partition is greater than the amount of available memory, perform one or more of the following tasks:
 - Add memory to the shared memory pool so that the shared memory pool has enough available memory to accommodate the I/O entitled memory required by the mobile partition. For instructions, see Changing the shared memory pool size by using the Integrated Virtualization Manager.
 - Remove one or more shared memory partitions from the shared memory pool until the shared memory pool has enough available memory to accommodate the I/O entitled memory required by the mobile partition. You can remove a logical partition from the shared memory pool by changing the memory mode of the logical partition from shared to dedicated. For instructions, see Managing memory properties for shared memory partitions.
 - If the amount of I/O entitled memory that is required by the mobile partition is equal to, or almost equal to, the amount of available memory, the shared memory pool is probably greatly overcommitted, which can affect performance. Consider adding more memory to the shared memory pool to reduce the degree to which the shared memory pool is overcommitted.

Attention: If you move an active logical partition whose I/O entitled memory mode is set to auto, the IVM does not automatically recalculate and reassign the I/O entitled memory for the mobile partition until you restart the mobile partition on the destination server. If you restart the mobile partition on the destination server and you plan to move the mobile partition back to the source server, you must verify that the shared memory pool on the source server has enough available memory to accommodate the new amount of I/O entitled memory required by the mobile partition.

Related information:

 Performance considerations for overcommitted shared memory partitions

Verifying that the destination server supports the virtual server network:

To move a mobile partition that uses the virtual server network (VSN), you must verify that the destination server also uses VSN by using the Hardware Management Console (HMC).

To verify that the destination server uses VSN, complete the following steps:

1. In the navigation pane, open **Systems Management** and click **Servers**.
2. Select the server in the work pane.
3. From the **Tasks** menu, click **Properties**.
4. Click the **Capabilities** tab.
 - If **Virtual Server Network Phase 2 Capable** is **True**, the destination server uses VSN.
 - If **Virtual Server Network Phase 2 Capable** is **False**, the destination server does not use VSN. To move the mobile partition to the destination server, disable VSN on the source server.

5. Click **OK**.

Determining the virtual Ethernet switch name and mode in the destination server:

Determine the name and mode of the virtual Ethernet switches in the destination server by using the Hardware Management Console (HMC).

To determine the name and mode of the virtual Ethernet switches, complete the following steps:

1. Determine the name and mode of the virtual Ethernet switches on the source server:
 - a. In the navigation pane, expand **Systems Management**, click **Servers** and select the source server on which the mobile partition is located.
 - b. In the **Tasks** menu, click **Configuration > Virtual resources > Virtual Network Management**.
 - c. Record the name and mode of each virtual Ethernet switch from the VSwitch area.
2. Determine the name and mode of the virtual Ethernet switches on the destination server:
 - a. In the navigation pane, expand **Systems Management**, click **Servers** and select the destination server to which you are moving the mobile partition.
 - b. In the **Tasks** menu, click **Configuration > Virtual resources > Virtual Network Management**.
 - c. Record the name and mode of each virtual Ethernet switch from the VSwitch area.

Compare the name and mode of the virtual Ethernet switches in the source server from step 1 with the name and mode of the virtual Ethernet switches in the destination server from step 2. The results of the comparison can be one of the following:

- If the name and mode are identical, the mobile partition can be successfully moved from the source server to the destination server.
- If the switch does not exist on the destination server, a switch with the same name and mode is automatically created in the destination server during the migration process.
- If a switch with the same name and different mode exists in the destination server, a warning message is displayed.

Related tasks:

“Resuming the suspended mobile partition with HMC” on page 86

You can resume a suspended Linux logical partition on the server by using the Hardware Management Console (HMC) Version 7.7.2.0, or later.

Determining available processors on the destination server:

You can determine the available processors on the destination server and allocate more processors, if necessary, by using the Integrated Virtualization Manager (IVM).

You must be a super administrator to perform this task.

To determine the available processors on the destination server using the using the IVM, complete the following steps:

1. Determine how many processors the mobile partition requires:
 - a. From the Partition Management menu, click **View/Modify Partition**. The View/Modify Partition panel is displayed.
 - b. Select the logical partition for which you want to view the properties.
 - c. From the Tasks menu, click **Properties**. The Partition Properties panel is displayed.
 - d. Click the **Processing** tab and record the minimum, maximum, and available processing units settings.
 - e. Click **OK**
2. Determine the processors available on the destination server:

- a. From the **Partition Management** menu, click **View/Modify System Properties**. The View/Modify System Properties panel is displayed.
 - b. Select the **Processing** tab.
 - c. Record the **Current processing units available**.
 - d. Click **Apply**.
3. Compare the values from steps 1 and 2.
 - If the destination server has enough available processors to support the mobile partition, then continue with “IVM-managed systems: Preparing the source and destination servers for partition mobility” on page 123.
 - If the destination server does not have enough available processors to support the mobile partition, use the IVM to dynamically remove the processors from the logical partition or you can remove processors from logical partitions on the destination server.

Preparing the source and destination management partitions for partition mobility

You need to verify that the source and destination management partitions are configured correctly so that you can successfully move the mobile partition from the source server to the destination server. This includes verifying the version of the Integrated Virtualization Manager (IVM) and activating the PowerVM hardware feature for IBM PowerLinux servers.

To prepare the source and destination management partitions for active or inactive partition mobility, complete the following tasks.

Table 59. Preparation tasks for the IVM

IVM planning tasks	Active mobility task	Inactive mobility task	Information resources
1. Ensure that the IVM that manages the source server and the IVM that manages the destination server meet the following version requirements: <ul style="list-style-type: none"> • If the source server, the destination server, or both servers are POWER7 processor-based servers, ensure that the IVM or IVMs that manage the servers are at version 2.1.2 with fix pack 22.1 and service pack 1, or later. • If the source server or the destination server is a POWER6 processor-based server, ensure that the IVM that manages that server is at version 2.1.2 with fix pack 22, or later. 	X	X	Updating the Integrated Virtualization Manager
2. Ensure that the PowerVM hardware feature for IBM PowerLinux servers is activated.	X	X	Entering the activation code for PowerVM for IBM PowerLinux with the Integrated Virtualization Manager
3. If the mobile partition uses shared memory, verify that the shared memory pool on the destination server contains a paging space device that satisfies the size requirements of the mobile partition.	X	X	“Verifying that the destination shared memory pool contains an available paging space device” on page 130

Related concepts:

“Integrated Virtualization Manager in a partition mobility environment” on page 117

Learn about the Integrated Virtualization Manager (IVM) and how you can use it to move an active or inactive logical partition from one server to another server.

Verifying that the destination shared memory pool contains an available paging space device:

You can verify that the shared memory pool on the destination server contains a paging space device that satisfies the size requirements of the mobile partition by using the Integrated Virtualization Manager (IVM).

To verify that the shared memory pool on the destination server contains a paging space device that satisfies the size requirements of the mobile partition, complete the following steps from the IVM:

1. Identify the size requirements of the mobile partition. The paging space device for the Linux logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) must be at least the size of the maximum logical memory of the shared memory partition. To view the maximum logical memory of the mobile partition, complete the following steps:
 - a. In the navigation pane, click **View/Modify Partitions** under **Partition Management**. The View/Modify Partitions page is displayed.
 - b. Select the mobile partition.
 - c. From the **Tasks** menu, click **Properties**. The Partition Properties page is displayed.
 - d. Click the **Memory** tab.
 - e. Note the maximum logical memory. This is the size requirement for the paging space device for the mobile partition.
2. View the paging space devices that are currently assigned to the shared memory pool on the destination server:
 - a. In the navigation pane, click **View/Modify Shared Memory Pool** under **Partition Management**. The View/Modify System Properties page is displayed.
 - b. Expand **Paging Space Devices - Advanced**.
 - c. Take note of the size of each paging space device that is not assigned to any shared memory partitions.
3. Identify the amount of available space in the paging storage pool:
 - a. In the navigation pane, click **View/Modify Virtual Storage** under **Virtual Storage Management**. The View/Modify Virtual Storage page is displayed.
 - b. Click the **Storage Pools** tab.
 - c. Select the paging storage pool.
 - d. From the **Tasks** menu, click **Properties**. The Storage Pool Properties page is displayed.
 - e. Note the available size of the paging storage pool.
4. Determine whether the shared memory pool on the destination server has a suitable paging space device for the mobile partition. The shared memory pool on the destination server has a suitable paging space device if one of the following situations is true:
 - The paging storage pool has enough space to meet the size requirements of the mobile partition (the result of step 3 less the result of step 1 is greater than or equal to zero). When you move the mobile partition to the destination server (active partition mobility) or when you activate the mobile partition on the destination server (inactive partition mobility), the IVM automatically creates a paging space device for the mobile partition.
 - The shared memory pool contains a paging space device that is not assigned to any shared memory partitions and meets the size requirements of the mobile partition.
5. If the shared memory pool on the destination server does not have a suitable paging space device, complete one of the following tasks:
 - Extend the size of the paging storage pool until there is enough space for the IVM to automatically create a paging space device for the mobile partition. For instructions, see Modifying storage pools using the Integrated Virtualization Manager.


- Add a paging space device that meets the size requirements of the mobile partition to the shared memory pool. For instructions, see Adding or removing paging space devices by using the Integrated Virtualization Manager.

Related concepts:

“Integrated Virtualization Manager in a partition mobility environment” on page 117

Learn about the Integrated Virtualization Manager (IVM) and how you can use it to move an active or inactive logical partition from one server to another server.

Related information:

 Paging space devices on systems that are managed by the Integrated Virtualization Manager

Preparing the mobile partition for partition mobility

You need to verify that the mobile partition is configured correctly so that you can successfully move it from the source server to the destination server by using the Integrated Virtualization Manager (IVM). This includes tasks such as satisfying adapter requirements and operating system requirements for partition mobility.

To prepare the mobile partition for active or inactive partition mobility, complete the following tasks.

Table 60. Preparation tasks for the mobile partition

Mobile partition planning tasks	Active mobility task	Inactive mobility task	Information resources
1. Ensure that the operating system running in the mobile partition is the Linux operating system.	X	X	
2. Ensure that the operating system is at one of the following levels: <ul style="list-style-type: none"> • Red Hat Enterprise Linux version 5 Update 5, or later • SUSE Linux Enterprise Server 10 Service Pack 3, or later • SUSE Linux Enterprise Server 11 Service Pack 1, or later <p>Earlier versions of the Linux operating system can participate in inactive partition mobility if the operating systems support virtual devices and POWER6 or POWER7 processor-based servers.</p>	X	X	
3. Ensure that the DynamicRM tool package is installed.	X		Service and productivity tools for Linux POWER servers
4. Ensure that the source and destination management partitions can communicate to each other.	X	X	
5. Verify that the processor compatibility mode of the mobile partition is supported on the destination server.	X	X	“Verifying the processor compatibility mode of the mobile partition” on page 133
6. Ensure that the mobile partition is not part of a partition workload group.	X	X	“Removing the mobile partition from a partition workload group” on page 134
7. Ensure that the mobile partition does not have physical I/O adapters. <p>Attention: During inactive migration, the IVM automatically removes any physical I/O adapters that are assigned to the mobile partition.</p>	X		Dynamically managing physical adapters

Table 60. Preparation tasks for the mobile partition (continued)

Mobile partition planning tasks	Active mobility task	Inactive mobility task	Information resources
8. Ensure that the mobile partition does not use Host Ethernet Adapters (or Integrated Virtual Ethernet).			Assigning Host Ethernet Adapter port to a logical partition
9. Ensure that the applications running in the mobile partition are mobility-safe or mobility-aware.	X		“Software applications that recognize partition mobility” on page 42

Configuring the Virtual I/O Server for the VSN capability:

If you are using the Hardware Management Console (HMC) Version 7 Release 7.7.0, or later, you can use Virtual Station Interface (VSI) profiles with virtual Ethernet adapters in logical partitions and assign the Virtual Ethernet Port Aggregator (VEPA) switching mode to virtual Ethernet switches.

When you use the Virtual Ethernet Bridge (VEB) switching mode in virtual Ethernet switches, the traffic between logical partitions is not visible to the external switches. However, when you use the VEPA switching mode, the traffic between logical partitions is visible to the external switches. This visibility helps you to use features such as security that are supported by the advanced switching technology. Automated VSI discovery and configuration with the external Ethernet bridges simplifies the switch configuration for the virtual interfaces that are created with logical partitions. The profile-based VSI management policy definition provides flexibility during configuration and maximizes the benefits of automation.

The configuration requirements on the Virtual I/O Server (VIOS) to use the VSN capability follow:


- At least one VIOS logical partition that is servicing the virtual switch must be active and must support the VEPA switching mode.
- The external switches that are connected to the shared Ethernet adapter must support the VEPA switching mode.
- The **lldpd** daemon must be running on the VIOS and must be managing the shared Ethernet adapter.
- From the VIOS command-line interface, run the **chdev** command to change the value of the *lldpsvc* attribute of the shared Ethernet adapter device to yes. The default value of the *lldpsvc* attribute is *no*. Run the **lldpsync** command to notify the change to the running **lldpd** daemon.

Note: The *lldpsvc* attribute must be set to the default value before you remove the shared Ethernet adapter. Otherwise, removal of the shared Ethernet adapter fails.

- For redundancy shared Ethernet adapter setup, the trunk adapters might be attached to a virtual switch that is set to the VEPA mode. In this case, attach the control channel adapters of the shared Ethernet adapter to another virtual switch that is always set to the virtual Ethernet bridging (VEB) mode. The shared Ethernet adapter that is in the high availability mode does not work when the control channel adapter that is associated with the virtual switches is in the VEPA mode.

Restriction: To use VSN capability, you cannot configure a shared Ethernet adapter to use link aggregation or an Etherchannel device as the physical adapter.

Related information:

 [chdev command](#)

Verifying the processor compatibility mode of the mobile partition:

You can use the Integrated Virtualization Manager (IVM) to determine whether the processor compatibility mode of the mobile partition is supported on the destination server, and update the mode, if necessary, so that you can successfully move the mobile partition to the destination server.

To verify that the processor compatibility mode of mobile partition is supported on the destination server using the IVM, complete the following steps:

1. Identify the processor compatibility modes that are supported by the destination server by entering the following command in the command line of the IVM on the destination server:

```
lssyscfg -r sys -F lpar_proc_compat_modes
```

Record these values so that you can refer to them later.

2. Identify the processor compatibility mode of the mobile partition on the source server:
 - a. From the Partition Management menu, click **View/Modify Partitions**. The View/Modify Partitions window is displayed.
 - b. In the work pane, select the mobile partition.
 - c. From the Tasks menu, select **Properties**. The Partition Properties window is displayed.
 - d. Select the **Processing** tab.
 - e. View the current and preferred processor compatibility modes for the mobile partition. Record these values so that you can refer to them later.

Restriction: If versions earlier than 2.1 of the IVM manage the source server, the IVM displays only the current processor compatibility mode for the mobile partition.

3. Verify that the processor compatibility mode that you identified in step 2 is on the list of supported processor compatibility modes that you identified in step 1 for the destination server. For active migrations, both the preferred and current processor compatibility modes of the mobile partition must be supported by the destination server. For inactive migrations, only the preferred processor compatibility mode must be supported by the destination server.

Attention: If the current processor compatibility mode of the mobile partition is the POWER5 mode, be aware that the POWER5 mode does not appear in the list of modes supported by the destination server. However, the destination server supports the POWER5 mode even though it does not appear in the list of supported modes.

4. If the preferred processor compatibility mode of the mobile partition is not supported by the destination server, use step 2 to change the preferred mode to a mode that is supported by the destination server. For example, the preferred mode of the mobile partition is the POWER7 mode and you plan to move the mobile partition to a POWER6 processor-based server. The POWER6 processor-based server does not support the POWER7 mode, but it does support the POWER6 mode. Therefore, you change the preferred mode to the POWER6 mode.
5. If the current processor compatibility mode of the mobile partition is not supported by the destination server, try the following solutions:
 - If the mobile partition is active, it is possible that the hypervisor has not had the opportunity to update the current mode of the mobile partition. Restart the mobile partition so that the hypervisor can evaluate the configuration and update the current mode of the mobile partition.
 - If the current mode of the mobile partition still does not match the list of supported modes that you identified for the destination server, use step 2 to change the preferred mode of the mobile partition to a mode that is supported by the destination server.

Then, restart the mobile partition so that the hypervisor can evaluate the configuration and update the current mode of the mobile partition.

For example, assume that the mobile partition runs on a POWER7 processor-based server and its current mode is the POWER7 mode. You want to move the mobile partition to a POWER6 processor-based server, which does not support the POWER7 mode. You change the preferred mode

of the mobile partition to the POWER6 mode and restart the mobile partition. The hypervisor evaluates the configuration and sets the current mode to the POWER6 mode, which is supported on the destination server.

Related concepts:

“Processor compatibility modes” on page 12

Processor compatibility modes enable you to move logical partitions between servers that have different processor types without upgrading the operating environments installed in the logical partitions.

“Processor compatibility modes” on page 99

Processor compatibility modes enable you to move logical partitions between servers that have different processor types without upgrading the operating environments installed in the logical partitions.

Removing the mobile partition from a partition workload group:

You can remove the mobile partition from a partition workload group by using the Integrated Virtualization Manager (IVM) so that you can move the mobile partition from the source server to the destination server.

A partition workload group identifies a set of logical partitions that are located on the same physical system. A partition workload group is defined when you use the IVM to configure a logical partition. The partition workload group is intended for applications that manage software groups. For a logical partition to participate in partition mobility, it cannot be assigned to a partition workload group.

To remove the mobile partition from a partition workload group using the IVM, complete the following steps:

1. From the Partition Management menu, click **View/Modify Partition**. The View/Modify Partition window is shown.
2. Select the logical partition that you want to remove from the partition workload group.
3. From the Tasks menu, and click **Properties**. The Partition Properties window is shown.
4. In the General tab, deselect **Partition workload group participant**.
5. Click **OK**.

Preparing the network configuration for partition mobility

You need to verify that the network configuration is configured correctly so that you can successfully move the mobile partition from the source server to the destination server by using the Integrated Virtualization Manager (IVM). This includes tasks such as configuring a virtual Ethernet bridge on the source and destination management partitions and creating at least one virtual Ethernet adapter on the mobile partition.

To prepare the network configuration for active or inactive partition mobility, complete the following tasks.

Note: Partition mobility fails if you have enabled one of the following security settings on the VIOS logical partitions:

- If you have set network security to the high mode by using the **viosecure** command on the VIOS command-line interface
- If you have enabled a profile that impacts network connectivity by using the **viosecure** command on the VIOS command-line interface

You can enable secure IP tunnels between the mover service partitions on the source and destination servers to perform partition mobility with these security settings. For more information, see “Configuring secure IP tunnels between the mover service partitions on the source and destination servers” on page 72.

Table 61. Preparation tasks for the network

Network planning tasks	Active mobility task	Inactive mobility task	Information resources
1. Configure a virtual Ethernet bridge on the source and destination management partitions using the IVM.	X	X	Configuring virtual Ethernet bridges on the managed system
2. Ensure that you connect the virtual Ethernet bridges on the source and destination management partitions to the network.	X	X	
3. Create at least one virtual Ethernet adapter on the mobile partition.	X		Creating a virtual Ethernet adapter
4. Activate the mobile partition to establish communication between the virtual Ethernet and management partition virtual Ethernet adapter.	X		Activating logical partitions
5. Verify that the operating system of the mobile partition recognizes the new Ethernet adapter.	X		Adapter management and configuration
6. Set up the LAN so that the mobile partition can continue to communicate with other necessary clients and servers after the migration is completed.	X	X	
7. Optional: Configure and enable secure IP tunnels between the mover service partitions on the source and destination servers.	X		“Configuring secure IP tunnels between the mover service partitions on the source and destination servers” on page 72
8. For VIOS partitions that are designated as mover service partitions, ensure that the network bandwidth between them is 1 GB or greater.	X		

Related concepts:

“Network configuration in a partition mobility environment” on page 119

In partition mobility that is managed by the Integrated Virtualization Manager (IVM), the network between the source and destination servers is used to pass the mobile partition state information and other configuration data from the source environment to the destination environment. The mobile partition uses the virtual LAN for network access.

Configuring secure IP tunnels between the mover service partitions on the source and destination servers:

With Virtual I/O Server (VIOS) 2.1.2.0, or later, you can configure secure IP tunnels between the mover service partitions on the source and destination servers. However, when both the source and destination servers are using the Virtual I/O Server 2.2.2.0, or later, the tunnels are created automatically depending on the security profile applied on the source VIOS.

Consider enabling secure IP tunnels between the mover service partition on the source server and the mover service partition on the destination server. For example, you might want to enable secure IP tunnels when the source and destination servers are not on a trusted network. Secure IP tunnels encrypt the partition state data that the mover service partition on the source server sends to the mover service partition on the destination server during active partition mobility.

Before you start, complete the following tasks:

1. Verify that the mover service partitions on the source and destination servers are at version 2.1.2.0, or later, by using the `ioslevel` command.
2. Obtain the IP address of the mover service partition on the source server.
3. Obtain the IP address of the mover service partition on the destination server.

4. Obtain the preshared authentication key for the source and destination mover service partitions.

To configure and enable secure IP tunnels, complete the following steps:

1. List the available secure tunnel agents by using the **lssvc** command. For example:

```
$lssvc  
ipsec_tunnel
```

2. List all the attributes that are associated with the secure tunnel agent by using the **cfgsvc** command. For example:

```
$cfgsvc ipsec_tunnel -ls  
local_ip  
remote_ip  
key
```

3. Configure a secure tunnel between the mover service partition on the source server and the mover service partition on the destination server by using the **cfgsvc** command:

```
cfgsvc ipsec_tunnel -attr local_ip=src_msp_ip remote_ip=dest_msp_ip key=key
```

where:

- *src_msp_ip* is the IP address of the mover service partition on the source server.
 - *dest_msp_ip* is the IP address of the mover service partition on the destination server.
 - *key* is the preshared authentication key for the mover service partitions on the source and destination servers. For example, abcderadf31231adsf.
4. Enable the secure tunnel by using the **startsvc** command. For example:

```
startsvc ipsec_tunnel
```

Note: When you apply the High, Payment Card Industry (PCI), or Department of Defence (DoD) security profiles, the secure tunnel is created and active partition mobility is performed over this secure channel. The secure channel that was created automatically gets destroyed when the partition mobility operation is complete.

Related concepts:

“Source and destination Virtual I/O Server logical partitions in a partition mobility environment” on page 31

Partition mobility that is managed by a Hardware Management Console (HMC) requires at least one Virtual I/O Server (VIOS) logical partition on the source server and at least one VIOS logical partition on the destination server.

“Integrated Virtualization Manager in a partition mobility environment” on page 117

Learn about the Integrated Virtualization Manager (IVM) and how you can use it to move an active or inactive logical partition from one server to another server.


“Network configuration in a partition mobility environment” on page 42

In partition mobility that is managed by the Hardware Management Console (HMC), the network between the source and destination servers is used to pass the mobile partition state information and other configuration data from the source environment to the destination environment. The mobile partition uses the virtual LAN for network access.

“Network configuration in a partition mobility environment” on page 119

In partition mobility that is managed by the Integrated Virtualization Manager (IVM), the network between the source and destination servers is used to pass the mobile partition state information and other configuration data from the source environment to the destination environment. The mobile partition uses the virtual LAN for network access.

Related information:

 [cfgsvc command](#)

 [startsvc command](#)

Preparing the virtual SCSI configuration for partition mobility

You need to verify that the virtual SCSI configuration is configured correctly so that you can successfully move the mobile partition from the source server to the destination server by using the Integrated Virtualization Manager (IVM). This includes tasks such as verifying the reserve_policy of the physical volumes, and verifying that the virtual devices have the same unique identifier, physical identifier, or IEEE volume attribute.

The destination server must provide the same virtual SCSI configuration as the source server. In this configuration, the mobile partition can access its physical storage on the storage area network (SAN) after it moves to the destination server.

To prepare the virtual SCSI configuration for active or inactive partition mobility, complete the following tasks.

Table 62. Preparation tasks for the virtual SCSI configuration on systems that are managed by the IVM

Storage planning tasks	Active mobility task	Inactive mobility task	Information resources
1. Verify that the physical storage that is used by the mobile partition is assigned to the management partition on the source server and to the management partition on the destination server.	X	X	IBM System Storage SAN Volume Controller
2. Verify that the reserve attributes on the physical volumes are the same for the source and destination VIOS partitions.	X	X	"Setting the reserve policy attributes of a device" on page 75
3. Verify that the virtual devices have the same unique identifier, physical identifier, or an IEEE volume attribute.	X	X	Identifying exportable disks
4. Optional: Specify a new name for one or more virtual target devices to use on the destination Virtual I/O Server (VIOS) partition.	X	X	"Specifying a new name for a virtual target device to use on a destination management partition" on page 139
5. Verify that the mobile partition has access to the physical storage on the SAN.	X	X	"Verifying that the mobile partition has access to its physical storage" on page 139

Related concepts:

"Storage configuration in a partition mobility environment" on page 120

Learn about the virtual SCSI and virtual Fibre Channel configuration required for partition mobility that is managed by the Integrated Virtualization Manager (IVM).

Setting the reserve policy attributes of a device:

In some configurations, you must consider the reservation policy of the device on the Virtual I/O Server (VIOS).

The following table explains the situations in which the reservation policy of the device on the VIOS is important for systems that are managed by the Hardware Management Console (HMC) and the Integrated Virtualization Manager (IVM).

Table 63. Situations where the reservation policy of a device is important

HMC-managed systems	IVM-managed systems
<ul style="list-style-type: none"> To use a Multipath I/O (MPIO) configuration at the client, none of the virtual Small Computer Serial Interface (SCSI) devices on the VIOS can reserve the virtual SCSI device. Set the <code>reserve_policy</code> attribute of the device to <code>no_reserve</code>. For virtual SCSI devices used with Live Partition Mobility or the Suspend/Resume feature, the <code>reserve</code> attribute on the physical storage that is used by the mobile partition can be set as follows: <ul style="list-style-type: none"> You can set the <code>reserve_policy</code> attribute to <code>no_reserve</code>. You can set the <code>reserve_policy</code> attribute to <code>pr_shared</code> when the following products are at the following versions: <ul style="list-style-type: none"> HMC Version 7 release 3.5.0, or later VIOS Version 2.1.2.0, or later The physical adapters support the SCSI-3 Persistent Reserves standard <p>The <code>reserve</code> attribute must be the same on the source and destination VIOS partitions for successful partition mobility.</p> For PowerVM Active Memory Sharing or Suspend/Resume features, the VIOS automatically sets the <code>reserve</code> attribute on the physical volume to <code>no_reserve</code>. The VIOS performs this action when you add a paging space device to the shared memory pool. 	<p>For virtual SCSI devices used with Live Partition Mobility, the <code>reserve</code> attribute on the physical storage that is used by the mobile partition can be set as follows:</p> <ul style="list-style-type: none"> You can set the <code>reserve_policy</code> attribute to <code>no_reserve</code>. You can set the <code>reserve_policy</code> attribute to <code>pr_shared</code> when the following products are at the following versions: <ul style="list-style-type: none"> IVM Version 2.1.2.0, or later The physical adapters support the SCSI-3 Persistent Reserves standard <p>The <code>reserve</code> attribute must be the same on the source and destination management partitions for successful partition mobility.</p>

- From a VIOS partition, list the disks (or paging space devices) to which the VIOS has access. Run the following command:

```
lsdev -type disk
```

- To determine the `reserve_policy` of a disk, run the following command, where `hdiskX` is the name of the disk that you identified in step 1 on page 75. For example, `hdisk5`.

```
lsdev -dev hdiskX -attr reserve_policy
```

The results might look like the following output:

```
..
reserve_policy no_reserve                Reserve Policy                True
```

Based on the information in Table 28 on page 75, you might need to change the `reserve_policy` so that you can use the disk in any of the described configurations.

- To set the `reserve_policy`, run the `chdev` command. For example:

```
chdev -dev hdiskX -attr reserve_policy=reservation
```

where:

- `hdiskX` is the name of the disk for which you want to set the `reserve_policy` attribute to `no_reserve`.
- `reservation` is either `no_reserve` or `pr_shared`.

- Repeat this procedure from the other VIOS partition.

Requirements:

- a. Although the `reserve_policy` attribute is an attribute of the device, each VIOS saves the value of the attribute. You must set the `reserve_policy` attribute from both VIOS partitions so that both VIOS partitions recognize the `reserve_policy` of the device.
- b. For partition mobility, the `reserve_policy` on the destination VIOS partition must be the same as the `reserve_policy` on the source VIOS partition. For example, if the `reserve_policy` on the source VIOS partition is `pr_shared`, the `reserve_policy` on the destination VIOS partition must also be `pr_shared`.
- c. With the `PR_exclusive` mode on SCSI-3 reserve, you cannot migrate from one system to another system.
- d. The `PR_key` value for the VSCSI disks on the source system and the target system must be different.

Verifying that the mobile partition has access to its physical storage:

You can use the Integrated Virtualization Manager (IVM) to verify that the mobile partition has access to its physical storage on the storage area network (SAN) so that the mobile partition can access its physical storage after it moves to the destination server.

For partition mobility to be successful, the mobile partition must have access to the same physical storage from both the source and destination environments. In the destination environment, the SAN host-attached adapter on the destination management partition must be connected to the same storage area network as the source management partition and have access to the same mobile partition physical storage as the source management partition.

To verify these connections using the IVM, complete the following steps:

1. From the Virtual Storage Management menu, click **View/Modify Virtual Storage**.
2. On the Virtual Disk tab, verify that the logical partition does not own any virtual disk.
3. On the Physical Volumes tab, verify the physical volumes mapped to the mobile partition are exportable. See *Identifying exportable disks* for more information.

If the information is incorrect, return to “Preparing the virtual SCSI configuration for partition mobility” on page 137 and complete the task associated with the incorrect information.

Specifying a new name for a virtual target device to use on a destination management partition:

Before you move a logical partition, you can specify a new name for a virtual target device, if needed. After you move the logical partition, the virtual target device assumes the new name on the Virtual I/O Server (VIOS) partition on the destination system.

Before you start, verify that the management partitions are at version 2.1.2.0, or later. This requirement applies to both the source management partition and the destination management partition.

Where possible, partition mobility preserves user-defined names of the virtual target devices on the destination system. Partition mobility does not preserve vtscsix IDs.

In some situations, partition mobility might not be able to preserve a user-defined name. For example, when the name is already in use on the destination VIOS partition.

If you want to maintain user-defined names on the destination VIOS partition, you can specify a new name for the virtual target device to use on the destination VIOS partition. If you do not specify a new name, partition mobility automatically assigns the next available vtscsix name to the virtual target device on the destination VIOS partition.

1. To view the names and mappings of the virtual target devices, run the **lsmap** command as follows. Run the command from the command-line interface on the source VIOS partition:

```
lsmap -all
```

The output might look like the following output:

SVSA	Physloc	Client Partition ID
vhost4	U8203.E4A.10D4431-V8-C14	0x0000000d
VTD	client3_hd0	
Status	Available	
LUN	0x8100000000000000	
Backing device	hdisk5	
Physloc	U789C.001.DQ1234#-P1-C1-T1-W500507630508C075-L4002402300000000	
VTD	client3_hd1	
Status	Available	
LUN	0x8200000000000000	
Backing device	hdisk6	
Physloc	U789C.001.DQ1234#-P1-C1-T1-W500507630508C075-L4002402400000000	

In this example, the user-defined names of the virtual target devices are `client3_hd0` and `client3_hd1`.

2. To specify a user-defined name for a virtual target device to use on the destination VIOS partition, run the **chdev** command as follows. Run the command from the command-line interface on the source VIOS partition:

```
chdev -dev dev_id -attr mig_name=partition_mobility_id
```

where:

- `dev_id` is the user-defined name of the virtual target device on the source VIOS partition.
- `partition_mobility_id` is the user-defined name that you want the virtual target device to have on the destination VIOS partition.

Related tasks:

“Validating the configuration for partition mobility” on page 142

You can use the Integrated Virtualization Manager (IVM) to validate the configuration of the source and destination systems for partition mobility. If the IVM detects a configuration or connection problem, it displays an error message with information to help you resolve the problem.

Preparing the virtual Fiber Channel configuration for partition mobility

You need to verify that the virtual Fiber Channel configuration is configured correctly so that you can successfully move the mobile partition from the source server to the destination server by using the Integrated Virtualization Manager (IVM). The verification includes tasks such as verifying the worldwide port names (WWPNs) of the virtual Fiber Channel adapters on the mobile partition, and verifying that the physical Fibre Channel adapters and the physical Fiber Channel switches support NPIV.

The destination server must provide the same virtual Fiber Channel configuration as the source server so that the mobile partition can access its physical storage on the storage area network (SAN) after it moves to the destination server.

To prepare the virtual Fiber Channel configuration for active or inactive partition mobility, complete the following tasks.

Table 64. Preparation tasks for the virtual Fiber Channel configuration on systems that are managed by the IVM

Storage planning tasks	Active mobility task	Inactive mobility task	Information resources
1. For each virtual Fiber Channel adapter on the mobile partition, verify that both (active and inactive) WWPNs for each virtual Fiber Channel adapter is zoned to the same storage port WWN on the switch or the storage area network (SAN).	X	X	<ul style="list-style-type: none"> To view the WWPNs that are assigned to a virtual Fiber Channel adapter, see Modifying partition properties IBM System Storage SAN Volume Controller
2. Verify that the physical Fiber Channel adapters that are assigned to the source and destination management partitions support NPIV. Run the lsnports command to view the physical ports on the physical Fiber Channel adapters that support NPIV.	X	X	Virtual I/O Server and Integrated Virtualization Manager commands
3. Verify that the switches to which the physical Fiber Channel adapters on both the source and destination management partitions are cabled support NPIV. Run the lsnports command to view the fabric support of the physical ports on the physical Fiber Channel adapters. If the fabric support is 1, the physical port is cabled to a switch that supports NPIV.	X	X	Virtual I/O Server and Integrated Virtualization Manager commands
4. Verify that the destination server provides enough available physical ports to support the virtual Fiber Channel configuration of the mobile partition.	X	X	“Verifying the number of physical Fibre Channel ports that are available on the destination management partition”

Related concepts:

“Storage configuration in a partition mobility environment” on page 120

Learn about the virtual SCSI and virtual Fibre Channel configuration required for partition mobility that is managed by the Integrated Virtualization Manager (IVM).

Related information:

 Redundancy configuration using virtual fibre channel adapters

Verifying the number of physical Fibre Channel ports that are available on the destination management partition:

You can use the Integrated Virtualization Manager (IVM) to verify that the management partition on the destination server provides a sufficient number of available physical ports for the mobile partition to maintain access to its physical storage on the storage area network (SAN) from the destination server.

To verify the number of physical ports that are available on the management partition on the destination server using the IVM, complete the following steps:

Tip: You can also use the **lslparmig** command to verify that the destination server provides enough available physical ports to support the virtual Fibre Channel configuration of the mobile partition

1. Determine the number of physical ports that the mobile partition uses on the source server:
 - a. From the **Partition Management** menu, click **View/Modify Partitions**. The View/Modify Partitions panel is displayed.
 - b. Select the mobile partition.
 - c. From the **Tasks** menu, click **Properties**. The Partition Properties panel is displayed.
 - d. Click the **Storage** tab.
 - e. Expand the **Virtual Fibre Channel** section

- f. Record the number of physical ports that are assigned to the mobile partition and click **OK**.
2. Determine the number of physical ports that are available on the management partition on the destination server:
 - a. From the **I/O Adapter Management** menu, click **View/Modify Virtual Fibre Channel**. The View/Modify Virtual Fibre Channel panel is displayed.
 - b. Record the number of physical ports with available connections.
3. Compare the information that you identified in step 1 on page 141 to the information that you identified in step 2.
 - If the number of physical ports with available connections from step 2 is greater than or equal to the number of physical ports that are assigned to the mobile partition from step 1 on page 141, the destination server provides enough available physical ports to support the mobile partition on the destination server.
 - If the number of physical ports with available connections from step 2 is less than the number of physical ports that are assigned to the mobile partition from step 1 on page 141, you need to add a physical Fibre Channel adapter (that supports N_Port ID Virtualization) to the destination server.

Related information:

 [Virtual I/O Server and Integrated Virtualization Manager commands](#)

Validating the configuration for partition mobility

You can use the Integrated Virtualization Manager (IVM) to validate the configuration of the source and destination systems for partition mobility. If the IVM detects a configuration or connection problem, it displays an error message with information to help you resolve the problem.

To validate the source and destination systems for partition mobility using the IVM, complete the following steps:

1. From the Partition Management menu, click **View/Modify Partitions**. The View/Modify Partitions panel is displayed.
2. Select the logical partition for which you want to migrate and from the Tasks menu, select **Migrate**.
3. Enter the **Remote IVM or HMC**, **Remote user ID**, and **Password** of the logical partition you plan to migrate.
4. Click **Validate** to confirm that the changed settings are acceptable for partition mobility.

Related concepts:

“Configuration validation for partition mobility” on page 96

You can learn about the tasks that the Integrated Virtualization Manager (IVM) performs to validate your system configuration for active and inactive partition mobility.

Related tasks:

“Specifying a new name for a virtual target device to use on a destination management partition” on page 139

Before you move a logical partition, you can specify a new name for a virtual target device, if needed.

After you move the logical partition, the virtual target device assumes the new name on the Virtual I/O Server (VIOS) partition on the destination system.

Moving the mobile partition

You can move an active or inactive logical partition from one server to another server by using the Integrated Virtualization Manager (IVM).

Before you move a logical partition from one server to another server, complete the following tasks from the IVM.

Table 65. Prerequisite tasks for moving a logical partition

Partition mobility prerequisite tasks	Active mobility task	Inactive mobility task	Information resources
1. Verify that you have completed all the required preparation tasks for partition mobility.	X	X	“Preparing for partition mobility” on page 123
2. Verify that the memory and processor resources are synchronized after dynamically adding or removing resources.	X	X	<ul style="list-style-type: none"> • Dynamically managing memory • Dynamically managing processing power
3. Verify that the source and destination servers are in the Operating state.	X	X	Viewing and modifying system properties
4. Verify that the mobile partition is powered off.		X	Modifying partition properties
5. Verify that the mobile partition is in the Operating state.	X		<ul style="list-style-type: none"> • Modifying partition properties • Activating a logical partition
6. Verify that the source and destination Virtual I/O Servers are active.	X	X	Activating a logical partition
7. Verify that all tape and CD jobs are completed or stopped.	X		
8. Run the migration validation tool on the IVM to verify that the servers, mobile partition, storage, and network are ready for partition mobility.	X	X	“Validating the configuration for partition mobility” on page 142

To move a logical partition from one server to another server by using the IVM, complete the following tasks:

1. From the Partition Management menu, click **View/Modify Partitions**. The View/Modify Partitions panel is displayed.
2. Select the logical partition that you want to move from the Tasks menu and select **Migrate**.
3. Enter the **Remote IVM**, **Remoter user ID**, and **Password** of the logical partition you plan to move.
4. Click **Migrate**.

After you move a logical partition from one server to another server, complete the following tasks.

Table 66. Postrequisite tasks for moving a logical partition

Partition mobility postrequisite tasks	Active mobility task	Inactive mobility task	Information resources
1. Activate the mobile partition on the destination server.		X	Activating a logical partition
2. Optional: Add dedicated I/O adapters to the mobile partition on the destination server	X	X	Dynamically managing physical adapters
3. If any virtual terminal connections were lost during the migration, re-establish the connections on the destination server.	X	X	Opening a virtual terminal session
4. Optional: Assign the mobile partition to a logical partition workload group.	X	X	Adding a client logical partition to the partition workload group
5. If mobility-unaware applications terminated on the mobile partition before its movement, restart those applications on the destination.	X		

Table 66. Postrequisite tasks for moving a logical partition (continued)

Partition mobility postrequisite tasks	Active mobility task	Inactive mobility task	Information resources
6. Optional: Back up the Virtual I/O Server management partition on the destination server to preserve the new virtual device mappings.	X	X	Backing up the Virtual I/O Server
7. Optional: Disable secure IP tunnels between the mover service partitions on the source and destination servers.	X		stopsvc command

Notices

This information was developed for products and services offered in the U.S.A. This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Dept. LRAS/Bldg. 903
11501 Burnet Road
Austin, TX 78758-3400
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, IBM License Agreement for Machine Code, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Programming interface information

This Live Partition Mobility publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM Virtual I/O Server Version 2.2.3.2.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

Terms and conditions

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights: Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the Publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER

EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.



Printed in USA