

Power Systems

Logical partitioning



Note

Before using this information and the product it supports, read the information in [“Notices” on page 237](#).

This edition applies to IBM® AIX® Version 7.2, to IBM AIX Version 7.1, to IBM AIX Version 6.1, to IBM i 7.4 (product number 5770-SS1), to IBM Virtual I/O Server Version 3.1.2 and to all subsequent releases and modifications until otherwise indicated in new editions. This version does not run on all reduced instruction set computer (RISC) models nor does it run on CISC models.

© **Copyright International Business Machines Corporation 2018, 2026.**

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

Contents

Logical partitioning.....	1
What's new in Logical partitioning.....	1
Logical partition overview.....	2
Hardware Management Console.....	2
Benefits of logical partitioning.....	7
Sharing resources between logical partitions.....	8
Managed systems.....	9
Manufacturing default configuration.....	10
Logical partitioning tools.....	10
Physical and virtual hardware resources.....	11
Application support for Linux logical partitions.....	60
Examples: Logically partitioned systems.....	60
Planning for logical partitions.....	61
Trusted Boot.....	63
Trusted Firewall.....	64
Preparing to configure Active Memory Expansion.....	64
Configuration requirements for shared memory.....	65
Configuration requirements and restrictions for the remote restart capability of a logical partition	67
Verifying that the server supports partitions that are capable of the simplified version of the remote restart feature.....	68
Verifying that the logical partition supports the simplified version of the remote restart feature....	69
Verifying that the server supports Virtual Trusted Platform Module	69
Enabling the platform keystore capability on a logical partition.....	69
Configuring the Virtual I/O Server for the VSN capability.....	70
Verifying that the server uses virtual server network.....	70
Verifying that the server supports single root I/O virtualization.....	71
Verifying the logical port limit and the owner of the SR-IOV adapter.....	71
Verifying that the server supports IBM i native I/O capability.....	72
Preparing to configure shared memory.....	72
Determining the size of the shared memory pool.....	74
Software licensing for IBM licensed programs on logical partitions.....	75
Minimum hardware configuration requirements for logical partitions.....	78
Partitioning with the HMC.....	79
Creating logical partitions.....	79
Configuring virtual resources for logical partitions.....	115
Managing the shared memory pool.....	125
Managing logical partitions.....	133
User mode access to the hardware accelerator.....	209
Performance.....	210
Security.....	220
Troubleshooting IBM i logical partitions.....	220
Troubleshooting the RMC connection between the logical partition and the HMC.....	224
Performance.....	226
Performance considerations for shared memory partitions.....	226
Adjusting the shared memory configuration to improve performance.....	233
Troubleshooting the RMC connection between the logical partition and the HMC.....	234
Notices.....	237
Accessibility features for IBM Power Systems servers.....	238

Privacy policy considerations	239
Programming interface information.....	240
Trademarks.....	240
Terms and conditions.....	240

Logical partitioning

You can set up, manage, and troubleshoot AIX, IBM i, Linux®, and Virtual I/O Server logical partitions by using the Hardware Management Console (HMC), or Virtual Partition Manager. By creating logical partitions, you can reduce the footprint of your data center by consolidating servers, and maximize the use of system resources by sharing resources across logical partitions.

What's new in Logical partitioning

Read about new or changed information in Logical partitioning since the previous update of this topic collection.

March 2021

The topic [“Dynamic Platform Optimizer”](#) on page 153 is updated with versions of IBM Power Systems servers on which the Dynamic Platform Optimizer (DPO) function is supported.

November 2020

The following topics are updated:

- [“Enabling the platform keystore capability on a logical partition”](#) on page 69
- [“Configuration requirements and restrictions for the remote restart capability of a logical partition”](#) on page 67
- [“Validating the simplified remote restart operation of a logical partition”](#) on page 85
- [“Remotely restarting a logical partition”](#) on page 86

May 2020

- The following topics are new or updated for single root I/O virtualization (SR-IOV) logical ports:
 - [“Configuration requirements and restrictions for the remote restart capability of a logical partition”](#) on page 67
 - [“Adding a single root I/O virtualization logical port to a logical partition dynamically”](#) on page 164
 - [“Viewing migratable SR-IOV logical ports and SR-IOV backup virtual devices”](#) on page 164
 - [“Modifying a single root I/O virtualization logical port that is assigned to a logical partition dynamically”](#) on page 164
 - [“Removing a single root I/O virtualization logical port from a logical partition dynamically”](#) on page 165
 - [“Creating a profile with migratable single root I/O virtualization logical ports”](#) on page 165
 - [“Recovering a migratable single root I/O virtualization logical port”](#) on page 165

October 2019

- The following topics are new or updated for single root I/O virtualization (SR-IOV) logical ports:
 - [“Configuration requirements and restrictions for the remote restart capability of a logical partition”](#) on page 67
 - [“Adding a single root I/O virtualization logical port to a logical partition dynamically”](#) on page 164
 - [“Viewing migratable SR-IOV logical ports and SR-IOV backup virtual devices”](#) on page 164
 - [“Modifying a single root I/O virtualization logical port that is assigned to a logical partition dynamically”](#) on page 164

- [“Removing a single root I/O virtualization logical port from a logical partition dynamically” on page 165](#)
- [“Creating a profile with migratable single root I/O virtualization logical ports” on page 165](#)
- [“Recovering a migratable single root I/O virtualization logical port” on page 165](#)

August 2018

- The following topic is new for the hardware accelerator:
 - [“Enabling user mode access to the hardware accelerator” on page 209](#)
- The following topic was updated for mapping the virtual I/O resources of a logical partition to the Virtual I/O Servers on the destination server:
 - [“Viewing the details of a simplified remote restart operation” on page 84](#)

Logical partition overview

Logical partitioning is the ability to make a server run as if it were two or more independent servers. When you logically partition a server, you divide the resources on the server into subsets called logical partitions. You can install software on a logical partition, and the logical partition runs as an independent logical server with the resources that you allocated to the logical partition.

You can assign processors, memory, and input/output devices to logical partitions. You can run AIX, IBM i, Linux, and the Virtual I/O Server in logical partitions. The Virtual I/O Server provides virtual I/O resources to other logical partitions with general-purpose operating systems.


Logical partitions share a few system attributes, such as the system serial number, system model, and processor feature code. All other system attributes can vary from one logical partition to another.

You can create a maximum of 1000 logical partitions on a server. You must use tools to create logical partitions on your servers. The tool that you use to create logical partitions on each server depends on the server model and the operating systems and features that you want to use on the server.

Hardware Management Console

The Hardware Management Console (HMC) is a hardware appliance that you can use to configure and control one or more managed systems. You can use the HMC to create and manage logical partitions and activate Capacity Upgrade on Demand. Using service applications, the HMC communicates with managed systems to detect, consolidate, and send information to service and support for analysis.

The HMC also provides terminal emulation for the logical partitions on your managed system. You can connect to logical partitions from the HMC itself, or you can set up the HMC so that you can connect to logical partitions remotely through the HMC. HMC terminal emulation provides a dependable connection that you can use if no other terminal device is connected or operational. HMC terminal emulation is useful during initial system setup before you configure your terminal of choice.



In this figure, you can see the logical partitions and the server firmware on the server. The server firmware is code that is stored in system flash memory on the server. The server firmware directly controls the resource allocations on the server and the communications between logical partitions on the server. The

HMC connects with the server firmware and specifies how the server firmware allocates resources to the managed system.

If you use a single HMC to manage a server, and the HMC malfunctions or becomes disconnected from the server firmware, then the server continues to run, but you cannot change the logical partition configuration of the server. If required, you can attach an extra HMC to act as a backup and to provide a redundant path between the server and service and support.

Partitioning by using the HMC is supported on all IBM Power Systems models, although some models require you to enter a PowerVM® Editions activation code before partitioning the managed system.

The PowerVM NovaLink architecture enables management of highly scalable cloud deployment by using the PowerVM technology and OpenStack solutions. The architecture provides a direct OpenStack connection to a PowerVM server. The NovaLink partition runs the Linux operating system and the partition runs on a server that is virtualized by PowerVM. The server is managed by PowerVC or other OpenStack solutions.

When a server is co-managed by the HMC and PowerVM NovaLink, and PowerVM NovaLink is in the master mode, you can run partition change operations only by using PowerVM NovaLink. If you want to run partition change operations by using the HMC, you must set the HMC to the master mode. Run the following command from the command line to set the HMC to the master mode:

```
chcomgmt -m <managed system> -o setmaster -t norm
```

Partition profile

A partition profile is a record on the Hardware Management Console (HMC) that specifies a possible configuration for a logical partition. When you activate a logical partition by using a partition profile, the managed system attempts to start the logical partition by using the configuration information in the partition profile.

A partition profile specifies the desired system resources for the logical partition and the minimum and maximum amounts of system resources that the logical partition can have. The system resources specified within a partition profile include processors, memory, and I/O resources. The partition profile can also specify certain operating settings for the logical partition. For example, you can set a partition profile such that, when the partition profile is activated, the logical partition is set to start automatically the next time that you power on the managed system.

Note: When a server is co-managed by the HMC, and PowerVM NovaLink, partition profiles are not supported.

Each logical partition on a managed system that is managed by an HMC has at least one partition profile. If you want, you can create more partition profiles with different resource specifications for your logical partition. If you create multiple partition profiles, you can designate any partition profile on the logical partition to be the default partition profile. The HMC activates the default profile if you do not select a specific partition profile to be activated. Only one partition profile can be active at one time. To activate another partition profile for a logical partition, you must shut down the logical partition before you activate the other partition profile.

A partition profile is identified by the logical partition ID and partition profile name. Logical partition IDs are whole numbers that are used to identify each logical partition that you create on a managed system, and partition profile names identify the partition profiles that you create for each logical partition. Each partition profile on a logical partition must have a unique partition profile name, but you can use a partition profile name for different logical partitions on a single managed system. For example, logical partition 1 cannot have more than one partition profile with a partition profile name of `normal`, but you can create a `normal` partition profile for each logical partition on the managed system.

When you create a partition profile, the HMC shows you all of the resources available on your system. The HMC does not verify whether another partition profile is currently using a portion of these resources. Therefore, it is possible for you to over commit resources. When you activate a logical partition by using a partition profile, the system attempts to start the logical partition by using the resources that are specified

in the partition profile. If the minimum resources that are specified in the partition profile are not available on the managed system, the logical partition cannot be started by using the partition profile.

For example, you have four processors on your managed system. Logical partition 1 with partition profile A has three processors, and logical partition 2 with partition profile B has two processors. If you attempt to activate both of these partition profiles at the same time, logical partition 2 with partition profile B fails to activate because you over committed processor resources.

When you shut down a logical partition and reactivate the logical partition by using a partition profile, the partition profile overlays the resource specifications of the logical partition with the resource specifications in the partition profile. Any resource changes that you made to the logical partition by using dynamic partitioning are lost when you reactivate the logical partition that uses a partition profile. This is preferred when you want to undo dynamic partitioning changes to the logical partition. However, this is not preferred if you want to reactivate the logical partition by using the resource specifications that the logical partition had when you shut down the managed system. Therefore, it is best to maintain your partition profiles up-to-date with the latest resource specifications. You can save the current configuration of the logical partition as a partition profile. Thereby, you do not need to change partition profiles manually.

If you shut down a logical partition whose partition profiles are not up-to-date, and if the logical partition is set to start automatically when the managed system starts, you can preserve the resource specifications on that logical partition by restarting the entire managed system by using the partition autostart power-on mode. When the logical partitions start automatically, the logical partitions have the resource specifications that the logical partitions had when you shut down the managed system.

You must activate a logical partition by activating a partition profile at least once. After that, you can activate the logical partition based on its current configuration data that is saved in the hypervisor. Logical partitions start faster when activated based on their current configuration data than when activated with a partition profile.

Processor resource assignment in partition profiles

When you create a partition profile for a logical partition, you can set up the allocated, minimum, and maximum amounts of processor resources that you want for the logical partition.

The allocated value is the resource amount that the logical partition gets if you do not over commit the resource on the managed system. If the allocated amount of resources is available when you activate the partition profile, then the logical partition starts with the allocated amount of resources. However, if the allocated amount of resources is not available when you activate the partition profile, then the resources on your managed system are over committed. If the amount of resources that are available on the managed system is equal to or greater than the minimum amount of resources in the partition profile, then the logical partition starts with the available amount of resources. If the minimum amount of resources is not met, then the logical partition does not start.

If the managed system allows the configuration of multiple shared processor pools, then you can limit the number of processors that are used by a specific group of logical partitions by configuring a shared processor pool for those logical partitions and reassigning those logical partitions to that shared processor pool. For example, if you use per-processor licensing for IBM i, and you have a limited number of IBM i licenses for your managed system, you can create a shared processor pool for the IBM i logical partitions on the managed system and set the maximum number of processing units for that shared processor pool to be equal to the number of IBM i licenses on the managed system. If you configure a shared processor pool and assign logical partitions to that shared processor pool, the number of processing units that are used by those logical partitions plus the number of processing units that are reserved for the use of uncapped logical partitions within the shared processor pool cannot exceed the maximum number of processing units that you set for that shared processor pool.

If you create a partition profile that is set to use shared processors, the HMC calculates a minimum, maximum, and allocated number of virtual processors for the partition profile. The calculation of virtual processors is based on the minimum, maximum, and allocated number of processing units that you specify for the partition profile. By default, the virtual processor settings are calculated as follows:

- The default minimum number of virtual processors is the minimum number of processing units (rounded up to the next whole number). For example, if the minimum number of processing units is 0.8, the default minimum number of virtual processors is 1.
- The default allocated number of virtual processors is the allocated number of processing units (rounded up to the next whole number). For example, if the allocated number of processing units is 2.8, the default allocated number of virtual processors is 3.
- The default maximum number of virtual processors is the maximum number of processing units rounded up to the next whole number and multiplied by two. For example, if the maximum number of processing units is 3.2, the default maximum number of virtual processors is 8 (four times 2).

When you activate the logical partition that uses the partition profile on the HMC, the allocated number of virtual processors is assigned to the logical partition. You can then use dynamic partitioning to change the number of virtual processors to any number between the minimum and maximum values, provided the number of virtual processors is greater than the number of processing units that are assigned to the logical partition. Before you change the default settings, performance modeling must be performed.

For example, you create a partition profile on the HMC with the following processor unit settings.

Minimum processing units 1.25
 Allocated processing units 3.80
 Maximum processing units 5.00

The default virtual processor settings for this partition profile on the HMC are as follows.

Minimum virtual processors 2
 Allocated virtual processors 4
 Maximum virtual processors 10

When you activate the logical partition by using this partition profile on the HMC, four processors are available to the operating system because the logical partition is activated with the allocated value of four virtual processors. Each of these virtual processors has 0.95 processing units to support the work that is assigned to the processor. After the logical partition is activated, you can use dynamic partitioning to change the number of virtual processors on the logical partition to any between number 2 - 10, provided the number of virtual processors is greater than the number of processing units that are assigned to the logical partition. If you increase the number of virtual processors, less processing power is available to support the work that is assigned to each processor.

Related concepts

Processors

A *processor* is a device that processes programmed instructions. The more processors that you assign to a logical partition, the greater the number of concurrent operations that the logical partition can run at any given time.

Memory resource assignment in partition profiles

When you create a partition profile for a logical partition, you can set up the allocated, minimum, and maximum amounts of memory resources that you want for the logical partition.

When you create a partition profile that is set to use dedicated memory, the allocated, minimum, and maximum amounts of memory that you specify refer to physical memory in the system. If the allocated amount of physical memory is available on the managed system when you activate the partition profile, the logical partition starts with the allocated amount of physical memory. However, if the allocated amount of physical memory is not available when you activate the partition profile, the physical memory on your managed system is over committed. In that case, if the amount of physical memory that is available on the managed system is equal to or greater than the minimum amount of physical memory in the partition profile, the logical partition starts with the available amount of physical memory. If the minimum amount of physical memory is not available, then the logical partition does not start.

When you create a partition profile that is set to use shared memory, the allocated, minimum, and maximum amounts of memory that you specify refer to logical memory. When you activate the partition profile, the logical partition starts with the allocated amount of logical memory. You can dynamically add

and remove logical memory to and from a running logical partition within the minimum and maximum values set in the partition profile.

Related concepts

Memory

Processors use memory to temporarily hold information. Memory requirements for logical partitions depend on the logical partition configuration, I/O resources assigned, and applications used.

I/O device assignment in partition profiles

I/O devices are assigned to partition profiles either on a slot-by-slot basis, or on logical port basis in the case of shared mode single root I/O virtualization (SR-IOV) adapters. For I/O devices that are assigned to partition profiles on a slot-by-slot basis, most I/O devices can be assigned to a partition profile on the HMC as required or as allocated. For SR-IOV logical ports, I/O devices are always assigned to a profile as required.

- If an I/O device is assigned to a partition profile as required, then the partition profile cannot be successfully activated if the I/O device is unavailable or is in use by another logical partition. Also, after the logical partition starts, you cannot use dynamic partitioning to remove the required I/O device from the running logical partition or move the required I/O device to another logical partition. This setting is suitable for devices that are required for the continuous operation of the logical partition (such as disk drives).
- If an I/O device is assigned to a partition profile as desired, then the partition profile can be successfully activated if the I/O device is unavailable or is in use by another logical partition. The desired I/O device can also be unconfigured in the operating system or system software and removed from the running logical partition or moved to another logical partition by using dynamic partitioning. This setting is suitable for devices that you want to share among multiple logical partitions (such as optical drives or tape drives).

The exception to this rule is host channel adapters (HCAs), which are added to partition profiles on the HMC as required. Each physical HCA contains a set of 64 globally unique IDs (GUIDs) that can be assigned to partition profiles. You can assign multiple GUIDs to each partition profile, but you can assign only one GUID from each physical HCA to each partition profile. Also, each GUID can be used by only one logical partition at a time. You can create multiple partition profiles with the same GUID, but only one of those partition profiles can be activated at a time.

You can change the required or desired setting within any partition profile for any I/O device at any time. Changes to the required or desired setting for an I/O device take effect immediately, even if the logical partition is running. For example, you want to move a tape device from one running logical partition to another, and the I/O device is required in the active partition profile for the source logical partition. You can access the active partition profile for the source logical partition, set the tape device to be allocated, and then unconfigure and move the tape device to the other logical partition without having to restart either logical partitions.

If you create an IBM i logical partition by using the HMC, you must tag I/O devices to perform certain functions for that IBM i logical partition.

Related concepts

Tagged resources for IBM i logical partitions

When you create an IBM i logical partition using the Hardware Management Console (HMC), you must tag I/O adapters (IOAs) to perform specific functions for the IBM i logical partition.

Related reference

Virtual adapters

With virtual adapters, you can connect logical partitions with each other without using physical hardware. Operating systems can display, configure, and use virtual adapters just like they can display, configure, and use physical adapters. Depending on the operating environment used by the logical partition, you can create virtual Ethernet adapters, virtual Fibre Channel adapters, virtual Small Computer Serial Interface (SCSI) adapters, and virtual serial adapters for a logical partition.

Partition profiles that use all of the system resources

You can create partition profiles on your HMC that specify all of the resources on the managed system. If you activate a logical partition by using such a partition profile, then the managed system assigns all of its resources to the logical partition.

If you add more resources to the managed system, the managed system automatically assigns the added resources to the logical partition when the profile is activated. The profile must be activated while the server is in the partition standby state because restarting the logical partition automatically does not assign newly added processor and memory resources. You do not have to change the partition profile for the managed system to assign the additional resources to the logical partition.

You cannot activate a logical partition that uses a partition profile that specifies all of the system resources if any other logical partition is running. However, after the logical partition is activated with all of the system resources, you can remove most processor and memory resources and all I/O resources from the logical partition by using dynamic partitioning. This allows you to start other logical partitions by using the resources that you remove from the logical partition. An implicit minimum amount of processor and memory resources is reserved for the logical partition that uses all of the system resources. Hence, you cannot remove all processor and memory resources from such a logical partition.

System profile

A system profile is an ordered list of partition profiles that is used by the Hardware Management Console (HMC) to start the logical partitions on a managed system in a specific configuration.

When you activate the system profile, the managed system attempts to activate each partition profile in the system profile in the order specified. A system profile helps you to activate or change the managed system from one complete set of logical partition configurations to another.

You can create a system profile and specify more resources to partition profiles than the resources that are available on the managed system. You can use the HMC to validate the system profile against the currently available system resources and against the total system resources. Validating your system profile ensures that your I/O devices and processing resources are not over committed, and it increases the likelihood that the system profile can be activated. The validation process estimates the amount of memory that is needed to activate all of the partition profiles in the system profile. A system profile can pass validation but might not have enough memory to be activated.

System profiles cannot include partition profiles that specify shared memory. In other words, logical partitions that use shared memory cannot be activated by using a system profile.

Benefits of logical partitioning

When you create logical partitions on your server, you can consolidate servers, share system resources, create mixed environments, and run integrated clusters.

The following scenarios illustrate the benefits of partitioning your server:

Consolidating servers

A logically partitioned server can reduce the number of servers that are required within an enterprise. You can consolidate several servers into a single logically partitioned system. This eliminates the need for, and expense of, more equipment.

Sharing resources

You can quickly and easily move hardware resources from one logical partition to another logical partition as needs change. Technologies such as the Micro-Partitioning® technology, allow for processor resources to be shared automatically among logical partitions that use a shared processor pool. Similarly, the PowerVM Active Memory Sharing technology allows for memory resources to be shared automatically among logical partitions that use the shared memory pool. Other technologies, such as dynamic partitioning, allow for resources to be manually moved to, from, and between running logical partitions without shutting down or restarting the logical partitions.

Maintaining independent servers

Dedicating a portion of the resources (disk storage unit, processors, memory, and I/O devices) to a logical partition achieves logical isolation of software. If configured correctly, logical partitions also

have some hardware fault tolerance. Batch and 5250 online transaction processing (OLTP) workloads, which might not run together on a single machine, can be isolated and run efficiently in separate partitions.

Creating a mixed production and test environment

You can create a combined production and test environment on the same server. The production logical partition can run your main business applications, and the test logical partition is used to test software. A failure in a test logical partition, while not necessarily planned, does not disrupt normal business operations.

Merging production and test environments

Partitioning enables separate logical partitions to be allocated for production and test servers, eliminating the need to purchase more hardware and software. When testing completes, the resources that are allocated to the test logical partition can be returned to the production logical partition or elsewhere. As new projects are developed, they can be built and tested on the same hardware on which they are eventually deployed.

Running integrated clusters

Using high-availability application software, your partitioned server can run as an integrated cluster. You can use an integrated cluster to protect your server from most unscheduled failures within a logical partition.

Although there are many benefits of creating logical partitions, consider the following points before you choose to use logical partitions:

- Processor and memory failures might result in the failure of the entire server with all of its logical partitions. (The failure of a single I/O device affects only the logical partition to which the I/O device belongs.) To reduce the possibility of system failure, you can use the Advanced System Management Interface (ASMI) to set the server to unconfigure failing processors or memory modules automatically. After the server unconfigures the failing processor or memory module, the server continues running without using the unconfigured processor or memory module.
- Administering a consolidated system might be more difficult in some ways than administering multiple smaller systems, particularly if the resources in the consolidated system are used at a level close to their capacity. If you anticipate that you will use your server at a level close to its capacity, consider ordering a server model that is capable of Capacity on Demand (CoD).

Related information

[Capacity on Demand](#)

Sharing resources between logical partitions

Although each logical partition acts as an independent server, the logical partitions on a server can share some types of resources with each other. The ability to share resources among many logical partitions allows you to increase resource utilization on the server and to move the server resources to where they are needed.

The following list illustrates some of the ways in which logical partitions can share resources. For some server models, the features that are mentioned in this list are options for which you must obtain and enter an activation code:

- The Micro-Partitioning technology (or shared processing) allows logical partitions to share the processors in shared processor pools. Each logical partition that uses shared processors is assigned a specific amount of processor power from its shared processor pool. By default, each logical partition is set such that the logical partition uses no more than its assigned processor power. Optionally, you can set a logical partition such that the logical partition can use processor power that is not being used by other logical partitions in its shared processor pool. If you set the logical partition such that it can use unused processor power, the amount of processor power that the logical partition can use is limited by the virtual processor settings of the logical partition and by the amount of unused processor power available in the shared processor pool that is used by the logical partition.
- Logical partitions can share the memory in the shared memory pool by using the PowerVM Active Memory Sharing technology (or shared memory). Instead of assigning a dedicated amount of physical

memory to each logical partition that uses shared memory (also referred to as *shared memory partitions*), the hypervisor constantly provides the physical memory from the shared memory pool to the shared memory partitions as needed. The hypervisor provides portions of the shared memory pool that are not currently being used by shared memory partitions to other shared memory partitions that need to use the memory. When a shared memory partition needs more memory than the current amount of unused memory in the shared memory pool, the hypervisor stores a portion of the memory that belongs to the shared memory partition in auxiliary storage. Access to the auxiliary storage is provided by a Virtual I/O Server logical partition. When the operating system attempts to access data that is located in the auxiliary storage, the hypervisor directs a Virtual I/O Server to retrieve the data from the auxiliary storage and write it to the shared memory pool, so that the operating system can access the data. The PowerVM Active Memory Sharing technology is available with the PowerVM Enterprise Edition hardware feature, which also includes the license for the Virtual I/O Server software. Only 512 byte block devices are supported for PowerVM Active Memory Sharing.

- Dynamic partitioning allows you to manually move resources to, from, and between running logical partitions without shutting down or restarting the logical partitions. This allows you to share devices that logical partitions use occasionally. For example, if the logical partitions on your server use an optical drive occasionally, you can assign a single optical drive to multiple logical partitions as a desired device. The optical drive would belong to only one logical partition at a time, but you can use dynamic partitioning to move the optical drive between logical partitions as needed. Dynamic partitioning is not supported on servers that are managed by using the Virtual Partition Manager.
 - Virtual I/O allows logical partitions to access and use I/O resources on other logical partitions. For example, virtual Ethernet allows you to create a virtual LAN that connects the logical partitions on your server to each other. If one of the logical partitions on the server has a physical Ethernet adapter that is connected to an external network, you can configure the operating system of that logical partition to connect the virtual LAN with the physical Ethernet adapter. This allows the logical partitions on the server to share a physical Ethernet connection to an external network.
 - A Host Ethernet Adapter (HEA), or Integrated Virtual Ethernet (IVE), allows multiple logical partitions on the same server to share a single physical Ethernet adapter. Unlike most other types of I/O devices, you can never assign the HEA itself to a logical partition. Instead, multiple logical partitions can connect directly to the HEA and use the HEA resources. The logical partitions can then access external networks through the HEA without using an Ethernet bridge on another logical partition.
- Note:** HEA is not supported on POWER9 processor-based server.
- The single root I/O virtualization (SR-IOV) specification defines extensions to the PCI Express (PCIe) specification. SR-IOV allows virtualization of the physical ports of an adapter. Hence, the ports can be shared by multiple partitions that are running simultaneously. For example, a single physical Ethernet port appears as several separate physical devices.

Related concepts

Shared processors

Shared processors are physical processors whose processing capacity is shared among multiple logical partitions. The ability to divide physical processors and share them among multiple logical partitions is known as the *Micro-Partitioning* technology.

Shared memory

You can configure your system such that multiple logical partitions share a pool of physical memory. A shared memory environment includes the shared memory pool, logical partitions that use the shared memory in the shared memory pool, logical memory, I/O entitled memory, at least one Virtual I/O Server logical partition, and paging space devices.

Managed systems

A managed system is a single physical server plus the resources that are connected to the physical server. The physical server and the connected resources are managed by the physical server as a single unit.

Connected resources can include expansion units, towers, and drawers, and storage area network (SAN) resources that are assigned to the server.

You can install a single operating system on a managed system and use the managed system as a single server. Alternately, you can use a partitioning tool, such as the Hardware Management Console (HMC) to create multiple logical partitions on the managed system. The partitioning tool manages the logical partitions on the managed system.

Manufacturing default configuration

The manufacturing default configuration is the initial single partition setup of the managed system as received from your service provider.

When your system is in the manufacturing default configuration, you can install an operating system on the managed system and use the managed system as a nonpartitioned server. In this state, you do not have to manage the system by using a Hardware Management Console (HMC).

If you choose to attach an HMC to a managed system that is in the manufacturing default configuration for reasons other than partitioning (such as to activate Capacity on Demand), all of the physical hardware resources on the system are automatically assigned to the logical partition. If you add new physical hardware resources to the managed system, the resources are automatically assigned to the logical partition. However, to use the newly added resources, you must dynamically add the resources to the logical partition or restart the logical partition. You do not have to make any partitioning changes on the server if you do not want to do so.

However, if you use the HMC to create, delete, change, copy, or activate any logical partitions or partition profiles on the managed system, the system is then in the partition mode. You must then use the HMC to manage the managed system. If the server has at least one IBM i logical partition, then you must also change the managed system properties on the HMC so that one of the IBM i logical partitions on the managed system is the service partition for the managed system. If a managed system is managed by using an HMC, and you want to return the managed system to a nonpartitioned state, or if you want to partition the managed system with the or the Virtual Partition Manager, then you must follow a special procedure to reset the server.

Managed systems that are partitioned by using the Virtual Partition Manager are not managed by an HMC. If a managed system is managed by using the Virtual Partition Manager, then you do not have to reset the server to return the managed system to a nonpartitioned state. Also, you do not have to reset the server if you want to start using the Virtual Partition Manager instead of using an HMC. To start using an HMC, back up the data on each logical partition, attach the HMC to the server, create the logical partitions, and restore the data to the storage assigned to each logical partition.

Related concepts

Logical partitioning tools

You must use tools to create logical partitions on your servers. The tool that you use to create logical partitions on each server depends upon the server model and the operating systems and features that you want to use on the server.

Logical partitioning tools

You must use tools to create logical partitions on your servers. The tool that you use to create logical partitions on each server depends upon the server model and the operating systems and features that you want to use on the server.

Virtual Partition Manager

The Virtual Partition Manager is a feature of IBM i that allows you to create and manage one host partition and up to four client logical partitions, running either Linux or IBM i. You can use the Virtual Partition

Manager to create logical partitions on a server that does not have a Hardware Management Console (HMC).

To use the Virtual Partition Manager, you must first install IBM i on a nonpartitioned server. After you install IBM i, you can initiate a console session on IBM i and use System Service Tools (SST) to create and configure IBM i or Linux logical partitions. IBM i controls the resource allocations of the logical partitions on the server.

When you use the Virtual Partition Manager to create logical partitions on a server, SST can be used to create and manage the logical partitions. IBM Navigator for i offers an improved interface and more features for these functions. The console session that you use to access SST can be initiated by using Operations Console LAN.

Related information

[Virtual Partition Manager: A Guide to Planning and Implementation](#)
[Creating IBM i Client Partitions Using Virtual Partition Manager](#)

Physical and virtual hardware resources

When you create logical partitions on a managed system, you can assign the physical resources on the managed system directly to logical partitions. You can also share hardware resources among logical partitions by virtualizing those hardware resources. The methods that are used to virtualize and share hardware resources depend on the type of resource that you are sharing.

Processors

A *processor* is a device that processes programmed instructions. The more processors that you assign to a logical partition, the greater the number of concurrent operations that the logical partition can run at any given time.

You can set a logical partition to use either processors that are dedicated to the logical partition or processors that are shared with other logical partitions. If a logical partition uses dedicated processors, then you must assign processors (in increments of whole numbers) to the logical partition. A logical partition that uses dedicated processors cannot use any processing capacity beyond the processors that are assigned to the logical partition.

By default, all physical processors that are not dedicated to specific logical partitions are grouped together in a *shared processor pool*. You can assign a specific amount of the processing capacity in this shared processor pool to each logical partition that uses shared processors. Some models allow you to use the HMC to configure multiple shared processor pools. These models have a *default shared processor pool* that contains all the processor resources that do not belong to logical partitions that use dedicated processors or logical partitions that use other shared processor pools. The other shared processor pools on these models can be configured with a maximum processing unit value and a reserved processing unit value. The maximum processing unit value limits the total number of processors that can be used by the logical partitions in the shared processor pool. The reserved processing unit value is the number of processing units that are reserved for the use of uncapped logical partitions within the shared processor pool.

You can set a logical partition that uses shared processors to use as little as 0.10 processing units, which is approximately a 10th of the processing capacity of a single processor. When the firmware is at level 7.6, or later, you can set a logical partition that uses shared processors to use as little as 0.05 processing units, which is approximately a 20th of the processing capacity of a single processor. You can specify the number of processing units to be used by a shared processor logical partition down to the 100th of a processing unit. Also, you can set a shared processor logical partition such that, if the logical partition requires more processing capacity than its assigned number of processing units, the logical partition can use processor resources that are not assigned to any logical partition or processor resources that are assigned to another logical partition but that are not being used by the other logical partition. (Some server models might require you to enter an activation code before you can create logical partitions that use shared processors.)

You can assign up to the entire processing capacity on the managed system to a single logical partition, if the operating system and server model supports doing so. You can configure your managed system such that it does not comply with the software license agreement for your managed system, but you will receive out-of-compliance messages if you operate the managed system in such a configuration.

Automatic redistribution of work when a processor fails

If the server firmware detects that a processor is about to fail, or if a processor fails when the processor is not in use, then the server firmware creates a serviceable event. The server firmware can also unconfigure the failing processor automatically, depending upon the type of failure and the unconfiguration policies that you set up using the Advanced System Management Interface (ASMI). You can also unconfigure a failing processor manually using the ASMI.

When the server firmware unconfigures a failing processor, and if unassigned or unlicensed processors are not available on the managed system, the processor unconfiguration can cause the logical partition to which the processor is assigned to shut down. To avoid shutting down mission-critical workloads when your server firmware unconfigures a failing processor, you can use the HMC to set partition availability priorities for the logical partitions on your managed system. A logical partition with a failing processor can acquire a replacement processor from one or more logical partitions with a lower partition-availability priority. The managed system can dynamically reduce the number of processors used by shared processor partitions with lower partition-availability priorities and use the freed processor resources to replace the failing processor. If this does not provide enough processor resources to replace the failing processor, the managed system can shut down logical partitions with lower partition-availability priorities and use those freed processor resources to replace the failing processor. The acquisition of a replacement processor allows the logical partition with the higher partition-availability priority to continue running after a processor failure.

A logical partition can use processors only from logical partitions with lower partition-availability priorities. If all of the logical partitions on your managed system have the same partition-availability priority, then a logical partition can replace a failed processor only if the managed system has unlicensed or unassigned processors.

By default, the partition-availability priority of Virtual I/O Server logical partitions with virtual SCSI adapters is set to 191. The partition availability priority of all other logical partitions is set to 127 by default.

Do not set the priority of Virtual I/O Server logical partitions to be lower than the priority of the logical partitions that use the resources on the Virtual I/O Server logical partition. Do not set the priority of IBM i logical partitions with virtual SCSI adapters to be lower than the priority of the logical partitions that use the resources on the IBM i logical partition. If the managed system shuts down a logical partition because of its partition availability priority, all logical partitions that use the resources on that logical partition are also shut down.

If a processor fails when the processor is in use, then the entire managed system shuts down. When a processor failure causes the entire managed system to shut down, the system unconfigures the processor and restarts. The managed system attempts to start the logical partitions that were running at the time of the processor failure with their minimum processor values, in partition-availability priority order, with the logical partition with the highest partition-availability priority being started first. If the managed system does not have enough processor resources to start all of the logical partitions with their minimum processor values, then the managed system starts as many logical partitions as it can with their minimum processor values. If there are any processor resources remaining after the managed system has started the logical partitions, then the managed system distributes any remaining processor resources to the running logical partitions in proportion to their desired processor values.

Related concepts

Software licensing for IBM licensed programs on logical partitions

If you use IBM licensed programs such as AIX and IBM i on a server with logical partitions, consider how many software licenses are required for your logical partition configuration. Careful consideration of your software might help minimize the number of software licenses that you must purchase.

Processor resource assignment in partition profiles

When you create a partition profile for a logical partition, you can set up the allocated, minimum, and maximum amounts of processor resources that you want for the logical partition.

Related information

[Setting deconfiguration policies](#)

[Deconfiguring hardware](#)

Dedicated processors

Dedicated processors are whole processors that are assigned to a single logical partition.

If you choose to assign dedicated processors to a logical partition, you must assign at least one processor to that logical partition. Likewise, if you choose to remove processor resources from a dedicated logical partition, you must remove at least one processor from the logical partition.

On systems that are managed by a Hardware Management Console (HMC), dedicated processors are assigned to logical partitions using partition profiles.

By default, a powered-off logical partition using dedicated processors provides processors to uncapped logical partitions that use shared processors. If the uncapped logical partition needs additional processor resources, the uncapped logical partition can use the idle processors that belong to the powered-off dedicated logical partition, if the total number of processors used by the uncapped logical partition does not exceed the virtual processors assigned to the uncapped logical partition, and if the use of these idle processors does not cause the shared processor pool to exceed its maximum processing units. When you power on the dedicated logical partition while the uncapped logical partition is using the processors, the activated logical partition regains all of its processing resources. If you use the HMC, you can prevent dedicated processors from being used in the shared processor pool by disabling this function in the partition properties panels.

You can also set the properties of a logical partition that uses dedicated processors such that, unused processing cycles on those dedicated processors can be made available to uncapped logical partitions while the dedicated processor logical partition is running. You can change the processor sharing mode of the dedicated processor logical partition at any time, without having to shut down and restart the logical partition.

Related concepts

[Partition profile](#)

A partition profile is a record on the Hardware Management Console (HMC) that specifies a possible configuration for a logical partition. When you activate a logical partition by using a partition profile, the managed system attempts to start the logical partition by using the configuration information in the partition profile.

Shared processors

Shared processors are physical processors whose processing capacity is shared among multiple logical partitions. The ability to divide physical processors and share them among multiple logical partitions is known as the *Micro-Partitioning* technology.

Note: For some models, the Micro-Partitioning technology is an option for which you must obtain and enter a PowerVM Editions activation code.

By default, all physical processors that are not dedicated to specific logical partitions are grouped together in a *shared processor pool*. You can assign a specific amount of the processing capacity in this shared processor pool to each logical partition that uses shared processors. Some models allow you to use the HMC to configure multiple shared processor pools. These models have a *default shared processor pool* that contains all the processors that do not belong to logical partitions that use dedicated processors or logical partitions that use other shared processor pools. The other shared processor pools on these models can be configured with a maximum processing unit value and a reserved processing unit value. The maximum processing unit value limits the total number of processing units that can be used by the logical partitions in the shared processor pool. The reserved processing unit value is the number of processing units that are reserved for the use of uncapped logical partitions within the shared processor pool.

You can assign partial processors to a logical partition that uses shared processors. *Processing units* are a unit of measure for shared processing power across one or more virtual processors. One shared processing unit on one virtual processor accomplishes approximately the same work as one dedicated processor.

The minimum number of processing units depends on the firmware level.

Table 1. Firmware level and processing units per virtual processor	
Firmware level	Minimum number of processing units per virtual processor
FW740, or earlier	0.10
FW760, or later	0.05

Some server models allow logical partitions to use only a portion of the total active processors on the managed system. Hence, the full processing capacity of the managed system cannot be assigned to logical partitions. This is true for server models with one or two processors, where a large portion of processor resources is used as overhead.

When the firmware is at level FW760, or later, overall server performance can be impacted when too many virtual processors are configured on the managed system. You can verify the number of configured virtual processors by using the **lshwres** command from the HMC command line. An example of the output after running the **lshwres** command follows:

```
lshwres -m sysname -r proc --level sys -F
curr_sys_virtual_procs,max_recommended_sys_virtual_procs
4,240
```

where:

- `curr_sys_virtual_procs` indicates the current number of configured virtual processors.
- `max_recommended_sys_virtual_procs` indicates the recommended maximum number of configured virtual processors.

It is suggested that the number of configured virtual processors must not exceed the maximum number so that server performance is not affected.

The maximum number of active virtual processors for a shared processor partition is limited by a number of factors. On Power 795, Power 870, Power 880, Power 870C, Power 880C and Power 980 model servers, the firmware has a limit of 128 active shared virtual processors per partition. On all other models of POWER7, Power 8, and POWER9, the firmware has a limit of 64 active shared virtual processors per partition.

Note: The limits on the number of active virtual processors for a shared processor partition is applicable for the firmware, but different operating systems and different operating system versions might impose limits lower than the firmware limits.

On HMC-managed systems, shared processors are assigned to logical partitions that use partition profiles.

Logical partitions that use shared processors can have a sharing mode of capped or uncapped. An *uncapped logical partition* is a logical partition that can use more processor power than its assigned processing capacity. The amount of processing capacity that an uncapped logical partition can use is limited only by the number of virtual processors assigned to the logical partition or the maximum processing unit that is allowed by the shared processor pool that the logical partition uses. In contrast, a *capped logical partition* is a logical partition that cannot use more processor power than its assigned processing units.

For example, logical partitions 2 and 3 are uncapped logical partitions, and logical partition 4 is a capped logical partition. Logical partitions 2 and 3 are each assigned 3.00 processing units and four virtual processors. Logical partition 2 currently uses only 1.00 of its 3.00 processing units, but logical partition 3 currently has a workload demand that requires 4.00 processing units. Because logical partition 3 is

uncapped and has four virtual processors, the server firmware automatically allows logical partition 3 to use 1.00 processing units from logical partition 2. This increases the processing power for logical partition 3 to 4.00 processing units. Soon afterward, logical partition 2 increases its workload demand to 3.00 processing units. The server firmware therefore automatically returns 1.00 processing units to logical partition 2 so that logical partition 2 can use its full, assigned processing capacity again. Logical partition 4 is assigned 2.00 processing units and three virtual processors, but currently has a workload demand that requires 3.00 processing units. Because logical partition 4 is capped, logical partition 4 cannot use any unused processing units from logical partitions 2 or 3. However, if the workload demand of logical partition 4 decreases below 2.00 processing units, logical partitions 2 and 3 might use any unused processing units from logical partition 4.

By default, logical partitions that use shared processors are capped logical partitions. You can set a logical partition to be an uncapped logical partition if you want the logical partition to use more processing power than its assigned amount.

Although an uncapped logical partition can use more processor power than its assigned processing capacity, the uncapped logical partition can never use more processing units than its assigned number of virtual processors. Also, the logical partitions that use a shared processor pool can never use more processing units than the maximum processing units configured for the shared processor pool.

If multiple uncapped logical partitions need more processor capacity at the same time, the server can distribute the unused processing capacity to all uncapped logical partitions. This distribution process is determined by the uncapped weight of each of the logical partitions.

Uncapped weight is a number in the range of 0 through 255 that you set for each uncapped logical partition in the shared processor pool. On the HMC, you can choose from any of the 256 possible uncapped weight values. By setting the uncapped weight (255 being the highest weight), any available unused capacity is distributed to contending logical partitions in proportion to the established value of the uncapped weight. The default uncapped weight value is 128. When you set the uncapped weight to 0, no unused capacity is distributed to the logical partition.

When the firmware is at level FW830, or earlier, uncapped weight is used only when more virtual processors consume unused resources than the available physical processors in the shared processor pool. If no contention exists for processor resources, the virtual processors are immediately distributed across the physical processors, independent of their uncapped weights. This can result in situations where the uncapped weights of the logical partitions do not exactly reflect the amount of unused capacity.

For example, logical partition 2 has one virtual processor and an uncapped weight of 100. Logical partition 3 also has one virtual processor, but an uncapped weight of 200. If logical partitions 2 and 3 both require more processing capacity, and there is not enough physical processor capacity to run both logical partitions, logical partition 3 receives two more processing units for every additional processing unit that logical partition 2 receives. If logical partitions 2 and 3 both require more processing capacity, and there is enough physical processor capacity to run both logical partitions, logical partition 2 and 3 receive an equal amount of unused capacity. In this situation, their uncapped weights are ignored.

When the firmware is at level FW840, or later, if multiple partitions are assigned to a shared processor pool, the uncapped weight is used as an indicator of how the processor resources must be distributed among the partitions in the shared processor pool with respect to the maximum amount of capacity that can be used by the shared processor pool. For example, logical partition 2 has one virtual processor and an uncapped weight of 100. Logical partition 3 also has one virtual processor, but an uncapped weight of 200. If logical partitions 2 and 3 both require more processing capacity, logical partition 3 receives two additional processing units for every additional processing unit that logical partition 2 receives.

The server distributes unused capacity among all of the uncapped shared processor partitions that are configured on the server, regardless of the shared processor pools to which they are assigned. For example, if you configure logical partition 1 to the default shared processor pool and you configure logical partitions 2 and 3 to a different shared processor pool, all three logical partitions compete for the same unused physical processor capacity in the server, even though they belong to different shared processor pools.

Related concepts

[Sharing resources between logical partitions](#)

Although each logical partition acts as an independent server, the logical partitions on a server can share some types of resources with each other. The ability to share resources among many logical partitions allows you to increase resource utilization on the server and to move the server resources to where they are needed.

Partition profile

A partition profile is a record on the Hardware Management Console (HMC) that specifies a possible configuration for a logical partition. When you activate a logical partition by using a partition profile, the managed system attempts to start the logical partition by using the configuration information in the partition profile.

Virtual processors

A *virtual processor* is a representation of a physical processor core to the operating system of a logical partition that uses shared processors.

When you install and run an operating system on a server that is not partitioned, the operating system calculates the number of operations that it can perform concurrently by counting the number of processors on the server. For example, if you install an operating system on a server that has eight processors, and each processor can perform two operations at a time, the operating system can perform 16 operations at a time. In the same way, when you install and run an operating system on a logical partition that uses dedicated processors, the operating system calculates the number of operations that it can perform concurrently by counting the number of dedicated processors that are assigned to the logical partition. In both cases, the operating system can easily calculate how many operations it can perform at a time by counting the whole number of processors that are available to it.

However, when you install and run an operating system on a logical partition that uses shared processors, the operating system cannot calculate a whole number of operations from the fractional number of processing units that are assigned to the logical partition. The server firmware must therefore represent the processing power available to the operating system as a whole number of processors. This allows the operating system to calculate the number of concurrent operations that it can perform. A *virtual processor* is a representation of a physical processor to the operating system of a logical partition that uses shared processors.

The server firmware distributes processing units evenly among the virtual processors assigned to a logical partition. For example, if a logical partition has 1.80 processing units and two virtual processors, each virtual processor has 0.90 processing units to support its workload.

You can assign only a limited number of processing units for each virtual processor. The minimum number of processing units for each virtual processor is 0.10 (or ten virtual processors for every processing unit). When the firmware is at level FW760, or later, the minimum number of processing units is further lowered to 0.05 (or 20 virtual processors for every processing unit). The maximum number of processing units that can be assigned to each virtual processor is always 1.00. This means that a logical partition cannot use more processing units than the number of virtual processors that it is assigned, even if the logical partition is uncapped.

A logical partition generally performs best if the number of virtual processors is close to the number of processing units available to the logical partition. This lets the operating system manage the workload on the logical partition effectively. In certain situations, you might be able to increase system performance slightly by increasing the number of virtual processors. If you increase the number of virtual processors, you increase the number of operations that can run concurrently. However, if you increase the number of virtual processors without increasing the number of processing units, the speed at which each operation runs will decrease. The operating system also cannot shift processing power between processes if the processing power is split between many virtual processors.

On HMC-managed systems, virtual processors are assigned to logical partitions using partition profiles.

Related concepts

Partition profile

A partition profile is a record on the Hardware Management Console (HMC) that specifies a possible configuration for a logical partition. When you activate a logical partition by using a partition profile,

the managed system attempts to start the logical partition by using the configuration information in the partition profile.

Software and firmware requirements for processing units

The minimum number of processing units of a logical partition depends on the firmware level and the version of the operating system that is running on the logical partition.

The following table lists the firmware levels and the operating system versions.

<i>Table 2. Software and firmware requirements for processing units.</i>				
Minimum number of processing units per virtual processor	Firmware level	IBM i	AIX	Linux
0.10	FW740, or earlier	All	All	All
0.05	FW760, or later	All	Version 7 with Technology Level 2 or at Version 6 with Technology Level 8, or later.	A Linux distribution that supports the lower processor entitlement of 0.05 processing units per virtual processor

Memory

Processors use memory to temporarily hold information. Memory requirements for logical partitions depend on the logical partition configuration, I/O resources assigned, and applications used.

Memory can be assigned in increments of 16 MB, 32 MB, 64 MB, 128 MB, and 256 MB. The default memory block size varies according to the amount of configurable memory in the system.

<i>Table 3. Default memory block size used for varying amounts of configurable memory</i>	
Amount of configurable memory	Default memory block size
Less than 4 GB	16 MB
Greater than 4 GB up to 8 GB	32 MB
Greater than 8 GB up to 16 GB	64 MB
Greater than 16 GB up to 32 GB	128 MB
Greater than 32 GB	256 MB

A logical partition can grow based on the amount of memory initially allocated to it. Memory is added and removed to and from logical partitions in units of logical memory blocks. For logical partitions that are initially sized less than 256 MB, the maximum size to which a logical partition can grow is 16 times its initial size (up to the assigned maximum memory of the logical partition). For logical partitions that are initially sized 256 MB or larger, the maximum size to which the logical partition can grow is 64 times its initial size (up to the assigned maximum memory of the logical partition). The smallest increment for adding or removing memory to or from a logical partition is 16 MB.

The memory block size can be changed by using the Logical Memory Block Size option in the Advanced System Management Interface (ASMI). The machine default value should only be changed under direction from your service provider. To change the memory block size, you must be a user with administrator authority, and you must shut down and restart the managed system for the change to take effect. If the minimum memory amount in any partition profile on the managed system is less than the new memory block size, you must also change the minimum memory amount in the partition profile.

Each logical partition has a hardware page table (HPT). The HPT ratio is the ratio of the HPT size to the maximum memory value for the logical partition. The HPT is allocated in the server firmware memory overhead for the logical partition, and the size of the HPT can affect the performance of the logical partition. The size of the HPT is determined by the following factors:

- The HPT ratio of 1/64 is the default value for IBM i logical partitions, and 1/128 for AIX and Linux logical partitions.

Note: You can override the default value by using the HMC command-line interface to change the value in the partition profile.

- The maximum memory values that you establish for the logical partition (dedicated or shared)

On systems that are managed by a Hardware Management Console, memory is assigned to logical partitions using partition profiles.

Related concepts

Memory resource assignment in partition profiles

When you create a partition profile for a logical partition, you can set up the allocated, minimum, and maximum amounts of memory resources that you want for the logical partition.

Dedicated memory

Dedicated memory is physical system memory that you assign to a logical partition that uses dedicated memory (hereafter referred to as a *dedicated memory partition*), and is reserved for use by the dedicated memory partition until you remove the memory from the dedicated memory partition or delete the dedicated memory partition.

Depending on the overall memory in your system and the maximum memory values you choose for each logical partition, the server firmware must have enough memory to perform logical partition tasks. The amount of memory required by the server firmware varies according to several factors. The following factors influence server firmware memory requirements:

- Number of dedicated memory partitions
- Partition environments of the dedicated memory partitions
- Number of physical and virtual I/O devices used by the dedicated memory partitions
- Maximum memory values assigned to the dedicated memory partitions

Note: Firmware level updates can also change the server firmware memory requirements. Larger memory block sizes can exaggerate the memory requirement change.

When selecting the maximum memory values for each dedicated memory partition, consider the following points:

- Maximum values affect the hardware page table (HPT) size for each dedicated memory partition
- The logical memory map size for each dedicated memory partition

If the server firmware detects that a memory module has failed or is about to fail, the server firmware creates a serviceable event. The server firmware can also unconfigure the failing memory module automatically, depending on the type of failure and the deconfiguration policies that you set up using the Advanced System Management Interface (ASMI). You can also unconfigure a failing memory module manually using the ASMI. If a memory module failure causes the entire managed system to shut down, the managed system restarts automatically if the managed system is in normal IPL mode. When the managed system restarts itself, or when you restart the managed system manually, the managed system attempts to start the dedicated memory partitions that were running at the time of the memory module failure with their minimum memory values. If the managed system does not have enough memory to start all of the dedicated memory partitions with their minimum memory values, the managed system starts as many dedicated memory partitions as it can with their minimum memory values. If any memory is remaining after the managed system has started as many dedicated memory partitions as it can, the managed system distributes the remaining memory resources to the running dedicated memory partitions in proportion to the required memory values.

Setting huge-page memory values for AIX dedicated memory partitions

Specify the number of 16 GB pages to allocate to an AIX huge-page memory pool.

About this task

On managed systems that support huge-page memory, you can use the Hardware Management Console (HMC) to set the value for the huge-page memory pool. You can also specify values for the number of huge pages to allocate to logical partitions.

Using huge pages can improve performance in specific environments that require a high degree of parallelism, such as in DB2® database. You can specify huge-page memory that can be used for the shared-memory buffer pools in DB2. For logically partitioned systems, you can specify the minimum, wanted, and maximum number of huge pages to assign to a logical partition when you create the logical partition or partition profile.

To set the huge-page memory values, the system must be in the powered-off state. The new value takes effect when you restart the system.

Calculating huge-page memory requirements for AIX dedicated memory partitions

Calculate the value for the number of pages to allocate to an AIX huge-page memory pool.

About this task

To use huge-page memory, you must ensure that your system has adequate memory resources to dedicate to the huge-page memory pool. The huge-page memory pool is a region of system memory that is mapped as 16 GB page segments and is managed separately from the base memory of the system. Before you can specify the value for huge-page memory, you must determine which applications you are running and what the huge-page requirements are for your applications.

Determining huge-page memory requirements for your application

The huge-page memory pool can be used to enhance performance for DB2 in AIX operating systems. To determine this value, calculate the amount of memory required for the shared buffer pool to support your DB2 applications. Refer to the DB2 recommendations for buffer pool memory for your particular application.

Note: The huge page memory allocation cannot be changed dynamically. When you change the number of huge pages on the server, the server must be rebooted. Changing the number of assigned huge pages for a logical partition requires that you restart the logical partition.

Considerations for calculating the huge-page values

The amount of huge-page memory that you can allocate is dependent on the following factors:

- Total amount of licensed memory for your server
- Amount of available memory after configured memory resources are accounted for
- Number of physical I/O connections to the server (each I/O connection requires memory for the I/O tables, which can be distributed among the physical memory regions and reduces the memory available for huge pages)
- Base memory configuration for logical partitions (huge pages are not calculated as part of the configured-partition memory allocation)
- The requirements that define a huge page, that is each huge page requires 16 GB of contiguous real memory and must start on a 16 GB memory boundary
- Huge pages cannot span processing units. Each processing unit requires 32 GB to ensure at least one 16 GB huge page when all of the other considerations previously listed are taken into account.



Attention: The server firmware reduces the huge-page pool size to satisfy some of these dependencies. When this occurs, error log entries are generated to indicate that the huge-page pool size was reduced. The error log reference code is B700 5300. The reference code details indicates hexadecimal values that indicate why the huge-page pool size could not be satisfied. The

following example shows the possible entries and how to interpret the additional words in these entries:

- word 3 = 0x0000000100000106: This means that the huge-page pool was reduced to satisfy the system hardware configuration
 - word 4 = number of user-configured huge pages
 - word 5 = number of huge pages that could be provided
- word 3 = 0x0000000100000105: This means that the huge-page pool was reduced to satisfy the memory configuration of logical partitions
 - word 4 = number of huge pages before logical partitions were created
 - word 5 = firmware calculated number of huge pages after satisfying logical partition memory requirements
 - word 6 = number of huge pages in the pool

Calculating huge-page memory values

To calculate the server memory requirements to support huge pages, use the following steps:

Procedure

1. Determine the amount of base system memory and round that figure to the next 16 GB value.
2. Determine the number of I/O connection loops on your system and multiply the number by 16 GB. This calculation is required because the server needs a memory table for each I/O connection, and a 16 GB huge page cannot be located where an I/O table exists.
3. Take the larger of the values determined in step 1 and step 2. This is your base memory value.
4. Determine the number of huge pages that is required for your AIX applications. To determine this value, use the guidelines provided by your application documentation and the [AIX Performance Management](#). Multiply the number of anticipated huge pages by 16 GB. Add this figure to the base figure determined in step 3. The resulting figure provides an estimate of the amount of licensed memory required to satisfy the logical partition and huge-page pool memory requirements for your system.

Shared memory

You can configure your system such that multiple logical partitions share a pool of physical memory. A shared memory environment includes the shared memory pool, logical partitions that use the shared memory in the shared memory pool, logical memory, I/O entitled memory, at least one Virtual I/O Server logical partition, and paging space devices.

Related concepts

Sharing resources between logical partitions

Although each logical partition acts as an independent server, the logical partitions on a server can share some types of resources with each other. The ability to share resources among many logical partitions allows you to increase resource utilization on the server and to move the server resources to where they are needed.

Overview of shared memory

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 configure to use shared memory, share the memory in the pool with other shared memory partitions.

For example, you create a shared memory pool with 16 GB of physical memory. You then create three logical partitions, configure them to use shared memory, and activate the shared memory partitions. Each shared memory partition can use the 16 GB that are in the shared memory pool.

The hypervisor determines the amount of memory that is allocated from the shared memory pool to each shared memory partition based on the workload and memory configuration of each shared memory partition. When allocating the physical memory to the shared memory partitions, the hypervisor ensures

that each shared memory partition can access only the memory that is allocated to the shared memory partition at any given time. A shared memory partition cannot access the physical memory that is allocated to another shared memory partition.

The amount of memory that you assign to the shared memory partitions can be greater than the amount of memory in the shared memory pool. For example, you can assign 12 GB to shared memory partition 1, 8 GB to shared memory partition 2, and 4 GB to shared memory partition 3. Together, the shared memory partitions use 24 GB of memory, but the shared memory pool has only 16 GB of memory. In this situation, the memory configuration is considered over committed.

Over committed memory configurations are possible because the hypervisor virtualizes and manages all of the memory for the shared memory partitions in the shared memory pool as follows:

1. When shared memory partitions are not actively using their memory pages, the hypervisor allocates those unused memory pages to shared memory partitions that currently need them. When the sum of the physical memory that is currently used by the shared memory partitions is less than or equal to the amount of memory in the shared memory pool, the memory configuration is *logically over committed*. In a logically over committed memory configuration, the shared memory pool has enough physical memory to contain the memory that is used by all shared memory partitions at one point in time. The hypervisor need not store any data in auxiliary storage.
2. When a shared memory partition requires more memory than the hypervisor can provide to it by allocating unused portions of the shared memory pool, the hypervisor stores some of the memory that belongs to a shared memory partition in the shared memory pool and stores the remainder of the memory that belongs to the shared memory partition in auxiliary storage. When the sum of the physical memory that is currently used by the shared memory partitions is greater than the amount of memory in the shared memory pool, the memory configuration is *physically over committed*. In a physically over committed memory configuration, the shared memory pool does not have enough physical memory to contain the memory that is used by all the shared memory partitions at one point in time. The hypervisor stores the difference of the physical and shared memory in the auxiliary storage. When the operating system attempts to access the data, the hypervisor might need to retrieve the data from auxiliary storage before the operating system can access it.

Because the memory that you assign to a shared memory partition might not always reside in the shared memory pool, the memory that you assign to a shared memory partition is *logical memory*. Logical memory is the address space assigned to a logical partition, that the operating system perceives as its main storage. For a shared memory partition, a subset of the logical memory is backed up by physical main storage (or physical memory from the shared memory pool) and the remaining logical memory is kept in auxiliary storage.

A Virtual I/O Server logical partition provides access to the auxiliary storage, or paging space devices, which are required for shared memory partitions in an over committed memory configuration. A *paging space device* is a physical or logical device that is used by a Virtual I/O Server to provide the paging space for a shared memory partition. The *paging space* is an area of nonvolatile storage that is used to hold portions of a shared memory partition's logical memory that does not reside in the shared memory pool. When the operating system that runs in a shared memory partition attempts to access data, and the data is located in the paging space device that is assigned to the shared memory partition, the hypervisor sends a request to a Virtual I/O Server to retrieve the data and write it to the shared memory pool so that the operating system can access it.

On systems that are managed by a Hardware Management Console (HMC), you can assign up to two Virtual I/O Server (VIOS) logical partitions to the shared memory pool at a time. When you assign two paging VIOS partitions to the shared memory pool, you can configure the paging space devices such that both paging VIOS partitions have access to the same paging space devices. When one paging VIOS partition becomes unavailable, the hypervisor sends a request to the other paging VIOS partition to retrieve the data on the paging space device.

You cannot configure paging VIOS partitions to use shared memory. Paging VIOS partitions do not use the memory in the shared memory pool. You assign paging VIOS partitions to the shared memory pool so that they can provide access to the paging space devices for the shared memory partitions that are assigned to the shared memory pool.

Driven by workload demands from the shared memory partitions, the hypervisor manages over committed memory configurations by continually performing the following tasks:

- Allocating portions of physical memory from the shared memory pool to the shared memory partitions as needed.
- Requesting a paging VIOS partition to read and write data between the shared memory pool and the paging space devices as needed.

The ability to share memory among multiple logical partitions is known as the PowerVM Active Memory Sharing technology. The PowerVM Active Memory Sharing technology is available with the PowerVM Enterprise Edition for which you must obtain and enter a PowerVM Editions activation code. Only 512 byte block devices are supported for PowerVM Active Memory Sharing.

Example: A shared memory configuration that is logically overcommitted

When the sum of the physical memory that is currently used by the shared memory partitions is less than or equal to the amount of memory in the shared memory pool, the memory configuration is *logically over committed*. In a logically over committed memory configuration, the shared memory pool has enough physical memory to contain the memory that is used by all shared memory partitions at one point in time.

The following figure shows a server with shared memory configuration that is logically over committed.

Shared memory partition 1 is assigned 12 GB of logical memory, Shared memory partition 2 is assigned 8 GB of logical memory, and Shared memory partition 3 is assigned 4 GB of logical memory. Together, the shared memory partitions are assigned 24 GB of logical memory, which is more than the 16.25 GB of logical memory allocated to the shared memory pool. Therefore, the memory configuration is over committed.

Shared memory partition 1 currently uses 8 GB of physical memory, Shared memory partition 2 currently uses 4 GB of physical memory, and Shared memory partition 3 currently uses 4 GB of physical memory. Together, the shared memory partitions currently use 16 GB of physical memory, which is equal to the amount of physical memory available to them in the shared memory pool. Therefore, the memory configuration is logically over committed. In other words, the shared memory pool contains enough physical memory for the hypervisor to allocate unused memory pages to shared memory partitions that need them. All of the memory currently used by the shared memory partitions resides in the shared memory pool.

Related concepts

Performance considerations for over committed shared memory partitions

Learn about how the degree to which the memory configuration of a logical partition that uses shared memory (also referred to as a *shared memory partition*) is over committed affects the performance of the shared memory partition. In general, the less over committed the memory configuration of a shared memory partition, the better its performance.

Example: A shared memory configuration that is physically overcommitted

When the sum of the physical memory that is currently used by the shared memory partitions is greater than the amount of memory in the shared memory pool, the memory configuration is *physically over committed*. In a physically over committed memory configuration, the shared memory pool does not have enough physical memory to contain the memory that is used by all the shared memory partitions at one point in time. The hypervisor stores the difference of the physical and shared memory in the auxiliary storage.

The following figure shows a server with shared memory configuration that is physically over committed.

Shared memory partition 1 is assigned 12 GB of logical memory, Shared memory partition 2 is assigned 8 GB of logical memory, and Shared memory partition 3 is assigned 4 GB of logical memory. Together, the shared memory partitions are assigned 24 GB of logical memory, which is more than the 16.25 GB of logical memory allocated to the shared memory pool. Therefore, the memory configuration is over committed.

Shared memory partition 1 currently uses 8 GB of physical memory, shared memory partition currently uses 5 GB of physical memory, and shared memory partition 3 currently uses 4 GB of physical memory. Together, the shared memory partitions currently use 17 GB of physical memory, which is greater than the amount of physical memory available to them in the shared memory pool. Therefore, the memory configuration is physically overcommitted. In other words, the shared memory pool does not contain enough physical memory for the hypervisor to satisfy the memory needs of all the shared memory partitions without storing some of the memory in the paging space devices. In this example, the difference of 1 GB is stored in the paging space device that is assigned to Shared memory partition 2. When Shared memory partition 2 needs to access data, the hypervisor might need to retrieve the data from the paging space device before the operating system can access it.

Related concepts

Performance considerations for over committed shared memory partitions


Learn about how the degree to which the memory configuration of a logical partition that uses shared memory (also referred to as a *shared memory partition*) is over committed affects the performance of the shared memory partition. In general, the less over committed the memory configuration of a shared memory partition, the better its performance.

Data flow for shared memory partitions

When the operating system that runs in a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) needs to access data, the data must reside in the shared memory pool. Systems with overcommitted memory configurations require the hypervisor and at least one Virtual I/O Server (VIOS) logical partition that is assigned to the shared memory pool (hereafter referred to as *paging VIOS partition*) to move data between the shared memory pool and the paging space devices as needed.

In a shared memory configuration that is physically overcommitted (where the sum of the logical memory that is currently used by all the shared memory partitions is greater than the amount of memory in the shared memory pool), the hypervisor stores some of the logical memory that belongs to a shared memory partition in the shared memory pool and some of the logical memory in a paging space device. For the operating system in a shared memory partition to access its memory, the memory must be in the shared memory pool. Thus, when the operating system needs to access data that is stored on the paging space device, the hypervisor works with a paging VIOS partition to move the data from the paging space device to the shared memory pool so that the operating system can access it.

The following figure shows the data flow for shared memory.



P9HAT003-1

Figure 3. The process of managing data in a shared memory configuration that is overcommitted

In general, the data flows as follows:

1. The operating system that runs in a shared memory partition attempts to access data.
 - If the data is in the shared memory pool, processing continues with step “7” on page 28.
 - If the data is not in the shared memory pool, a page fault occurs. The hypervisor inspects the page fault and discovers that the hypervisor moved the data to the paging space device, thereby causing the page fault. Processing continues with step “2” on page 28. (If the operating system that runs in the shared memory partition moved the data to auxiliary storage, thereby causing the page fault, then the operating system must retrieve the data.)
2. The hypervisor sends a request to a paging VIOS partition to retrieve the data from the paging space device and to write it to the shared memory pool.
3. The paging VIOS partition searches the paging space device that is assigned to the shared memory partition and finds the data.
4. The paging VIOS partition writes the data to the shared memory pool.
5. The paging VIOS partition notifies the hypervisor that the data is in the shared memory pool.
6. The hypervisor notifies the operating system that it can access the data.
7. The operating system accesses the data in the shared memory pool.

Related concepts

Logical memory

Logical memory is the address space, assigned to a logical partition, that the operating system perceives as its main storage. For a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*), a subset of the logical memory is backed up by physical main storage and the remaining logical memory is kept in auxiliary storage.

Paging space device

You can learn about how the Hardware Management Console (HMC) allocates and manipulates paging space devices on systems that use shared memory.

Shared memory distribution

The hypervisor uses the memory weight of each logical partition that uses shared memory (hereafter referred to as *shared memory partitions*) to help determine which logical partitions receive more physical memory from the shared memory pool. To help optimize performance and memory use, the operating systems that run in shared memory partitions provide the hypervisor with information about how the operating system uses its memory to help the hypervisor determine which pages to store in the shared memory pool and which pages to store in the paging space devices.

Logical memory

Logical memory is the address space, assigned to a logical partition, that the operating system perceives as its main storage. For a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*), a subset of the logical memory is backed up by physical main storage and the remaining logical memory is kept in auxiliary storage.

You can configure minimum, maximum, desired, and assigned logical memory sizes for a shared memory partition.

Table 4. Logical memory sizes	
Logical memory size	Description
Minimum	The minimum amount of logical memory with which you want the shared memory partition to operate. You can dynamically remove logical memory from the shared memory partition down to this value.
Maximum	The maximum amount of logical memory that the shared memory partition is allowed to use. You can dynamically add logical memory to the shared memory partition up to this value.

Table 4. Logical memory sizes (continued)	
Logical memory size	Description
Desired	The amount of logical memory with which you want the shared memory partition to activate.
Assigned	The amount of logical memory that the shared memory partition can use. A shared memory partition does not have to use all of its assigned logical memory at any given time.

On systems that are managed by a Hardware Management Console (HMC), you configure the minimum, maximum, and desired logical memory sizes in the partition profile. When you activate the shared memory partition, the HMC assigns the desired logical memory to the shared memory partition.

The following figure shows a shared memory partition with its logical memory.




Figure 4. A shared memory partition that is assigned more logical memory than the amount of physical memory currently allocated to it

The figure shows a shared memory partition that is assigned 2.5 GB of logical memory. Its maximum logical memory is 3 GB and its minimum logical memory is 1 GB. You can change the assigned logical memory by dynamically adding or removing logical memory to or from the shared memory partition. You can dynamically add logical memory to the shared memory partition up to the maximum logical memory

size, and you can dynamically remove logical memory from the shared memory partition down to its minimum logical memory size.

The figure also shows that the amount of physical memory that is currently allocated to the shared memory partition from the shared memory pool is 2.1 GB. If the workload that runs in the shared memory partition currently uses 2.1 GB of memory and requires an additional 0.2 GB of memory, and the shared memory pool is logically overcommitted, the hypervisor allocates an additional 0.2 GB of physical memory to the shared memory partition by assigning memory pages that are not currently in use by other shared memory partitions. If the shared memory pool is physically overcommitted, the hypervisor stores 0.2 GB of the shared memory partition's memory in a paging space device. When the shared memory partition needs to access the data that resides in the paging space device, the hypervisor retrieves the data for the operating system.

The amount of physical memory allocated to the shared memory partition can be less than the minimum logical memory size. This is because the minimum logical memory size is a boundary for logical memory, not for physical memory. In addition to the minimum logical memory size, the maximum, desired, and assigned logical memory sizes also do not control the amount of physical memory assigned to the shared memory partition. Likewise, dynamically adding or removing logical memory to or from a shared memory partition does not change the amount of physical memory allocated to the shared memory partition. When you set the logical memory sizes and dynamically add or remove logical memory, you set or change the amount of memory that the operating system can use, and the hypervisor decides how to distribute that memory between the shared memory pool and the paging space device.

Related concepts

Data flow for shared memory partitions

When the operating system that runs in a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) needs to access data, the data must reside in the shared memory pool. Systems with overcommitted memory configurations require the hypervisor and at least one Virtual I/O Server (VIOS) logical partition that is assigned to the shared memory pool (hereafter referred to as *paging VIOS partition*) to move data between the shared memory pool and the paging space devices as needed.

Paging space device

You can learn about how the Hardware Management Console (HMC) allocates and manipulates paging space devices on systems that use shared memory.

Shared memory distribution

The hypervisor uses the memory weight of each logical partition that uses shared memory (hereafter referred to as *shared memory partitions*) to help determine which logical partitions receive more physical memory from the shared memory pool. To help optimize performance and memory use, the operating systems that run in shared memory partitions provide the hypervisor with information about how the operating system uses its memory to help the hypervisor determine which pages to store in the shared memory pool and which pages to store in the paging space devices.

Partition profile

A partition profile is a record on the Hardware Management Console (HMC) that specifies a possible configuration for a logical partition. When you activate a logical partition by using a partition profile, the managed system attempts to start the logical partition by using the configuration information in the partition profile.

Related tasks

Preparing to configure shared memory

Before you configure the shared memory pool and create logical partitions that use shared memory (hereafter referred to as *shared memory partitions*), you need to plan for the shared memory pool,

the shared memory partitions, the paging space devices, and the Virtual I/O Server logical partitions (hereafter referred to as *paging VIOS partitions*).

I/O entitled memory

I/O entitled memory is the maximum amount of physical memory (from the shared memory pool) that is guaranteed to be available to a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) for its I/O devices at any given time.

Each shared memory partition is entitled to some portion of the shared memory pool so that the I/O devices that are assigned to the shared memory partition have access to physical memory during I/O operations. If the minimum amount of memory that I/O devices require for I/O operations does not reside in the shared memory pool for the duration the device needs the memory, the device fails. Virtual adapters that are entitled to physical memory from the shared memory pool include virtual SCSI adapters, virtual Ethernet adapters, and virtual Fibre Channel adapters. Virtual serial adapters are not entitled to physical memory from the shared memory pool.

The following figure shows a shared memory partition with I/O entitled memory.




Figure 5. A shared memory partition whose I/O entitled memory is greater than the amount of physical memory that it currently uses for its I/O devices

The figure shows a shared memory partition with 128 MB of I/O entitled memory. The shared memory partition uses 64 MB of physical memory for its I/O devices, which is less than its I/O entitled memory of 128 MB.

As depicted in the previous figure, a shared memory partition might not use all of its I/O entitled memory at any given time. Unused portions of the I/O entitled memory assigned to a shared memory partition are available to the hypervisor to allocate to other shared memory partitions, if necessary. The hypervisor does not reserve unused portions of I/O entitled memory for the shared memory partition to use in the future. However, the hypervisor guarantees that the shared memory partition can use the entire portion of the I/O entitled memory that is assigned to it as needed. If the shared memory partition later requires some of its unused I/O entitled memory, the hypervisor must allocate enough physical memory from the shared memory pool to satisfy the new I/O memory requirement, without exceeding the I/O entitled memory that is assigned to the shared memory partition.

For example, you assign 128 MB of I/O entitled memory to a shared memory partition. The shared memory partition uses only 64 MB for its I/O devices. Thus, the hypervisor allocates 64 MB of physical memory from the shared memory pool to the shared memory partition for its I/O devices. The remaining 64 MB is available to the hypervisor to allocate to other shared memory partitions, if necessary. Later, you add two virtual adapters to the shared memory partition, each requiring 16 MB of memory. Thus, the shared memory partition needs an additional 32 MB of physical memory for its I/O devices. Because the shared memory partition currently uses only 64 MB of physical memory for its I/O devices and the shared memory partition is entitled to use up to 128 MB for its I/O devices, the hypervisor allocates an additional 32 MB of physical memory from the shared memory pool to the shared memory partition to accommodate the new virtual adapters. The shared memory partition now uses 96 MB of physical memory from the shared memory pool for its I/O devices.

Because unused portions of I/O entitled memory are available to the hypervisor to allocate elsewhere, for the amount of total physical memory that the hypervisor allocates from the shared memory pool to a shared memory partition can be less than the I/O entitled memory of the shared memory partition. The following figure shows this situation.




Figure 6. A shared memory partition whose I/O entitled memory is greater than the total amount of physical memory allocated to it

The figure shows a shared memory partition with 128 MB of I/O entitled memory. The shared memory partition uses 64 MB of physical memory for its I/O devices. The unused portion of the I/O entitled memory, 64 MB, is available to the hypervisor to allocate to other shared memory partitions, if necessary. The hypervisor allocates a total of 96 MB of physical memory from the shared memory pool to the shared memory partition, which is less than the I/O entitled memory of 128 MB.

When you create a shared memory partition, the Hardware Management Console (HMC) automatically sets the I/O entitled memory for the shared memory partition. When you activate a shared memory partition, the HMC sets the I/O entitled memory mode to the *auto* mode. In the auto mode, the HMC automatically adjusts the I/O entitled memory for the shared memory partition when you add or remove virtual adapters.

The I/O entitled memory mode can also be set to the *manual* mode. You can dynamically change the I/O entitled memory mode to the manual mode and then dynamically change the I/O entitled memory for the shared memory partition. When you add or remove a virtual adapter to or from the shared memory partition in manual mode, the HMC does not automatically adjust the I/O entitled memory. Therefore, you might need to dynamically adjust the I/O entitled memory when you dynamically add or remove adapters to or from the shared memory partition. On HMC-managed systems, you use the graphical interface to dynamically change the I/O entitled memory mode. When the I/O entitled memory mode is in the manual mode, you can also use the graphical interface to dynamically change the amount of I/O entitled memory that is assigned to a shared memory partition. When the I/O entitled memory mode is in the

manual mode, you can also use the **chhwres** command to dynamically change the amount of I/O entitled memory that is assigned to a shared memory partition. When you restart a shared memory partition, the I/O entitled memory mode is set to the auto mode regardless of what the I/O entitled memory mode was set to before you restarted the shared memory partition.

When the amount of physical memory that a shared memory partition uses for its I/O devices is equal to the I/O entitled memory that is assigned to the shared memory partition, the shared memory partition cannot use any more physical memory for its I/O devices. In this situation, the following actions can occur:

- The operating system that runs in the shared memory partition manages the I/O operations so that the workload that runs in the shared memory partition operates within the I/O entitled memory that is assigned to the shared memory partition. If the workload attempts to use more physical memory for I/O operations than the I/O entitled memory that is assigned to the shared memory partition, the operating system delays some I/O operations while it runs other I/O operations. In this situation, the I/O entitled memory of the shared memory partition constrains the I/O configuration of the shared memory partition because the operating system does not have enough physical memory to run all of the I/O operations simultaneously.
- When you dynamically add a virtual adapter to the shared memory partition and the I/O entitled memory mode is in the manual mode, the I/O configuration of the shared memory partition might become constrained, or the adapter might fail when you attempt to configure it. If the adapter fails, enough I/O entitled memory is not assigned to the shared memory partition to accommodate the new adapter. To resolve the problem, you can dynamically increase the amount of I/O entitled memory that is assigned to the shared memory partition, or you can remove some existing virtual adapters from the shared memory partition. When you remove virtual adapters from the shared memory partition, the physical memory that those adapters were using becomes available for the new adapter.
- When you dynamically add a virtual adapter to the shared memory partition and the I/O entitled memory mode is in the auto mode, the HMC automatically increases the I/O entitled memory assigned to the shared memory partition to accommodate the new adapter. If the HMC cannot increase the I/O entitled memory of the shared memory partition, enough physical memory is not available in the shared memory pool for the hypervisor to allocate to the shared memory partition and the adapter cannot be assigned to the shared memory partition. To resolve the problem, you can add physical memory to the shared memory pool, or you can remove some existing virtual adapters from the shared memory partition. When you remove virtual adapters from the shared memory partition, the physical memory that those adapters were using becomes available for the new adapter.

Paging VIOS partition

A Virtual I/O Server (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 are assigned to the shared memory pool (hereafter referred to as *shared memory partitions*).

When the operating system that runs in a shared memory partition attempts to access data, and the data is located in the paging space device that is assigned to the shared memory partition, the hypervisor sends a request to a paging VIOS partition to retrieve the data and write it to the shared memory pool so that the operating system can access it.

A paging VIOS partition is not a shared memory partition and does not use the memory in the shared memory pool. A paging VIOS partition provides access to the paging space devices for the shared memory partitions.

HMC


On systems that are managed by a Hardware Management Console (HMC), you can assign one or two paging VIOS partitions to the shared memory pool. When you assign a single paging VIOS partition to the shared memory pool, the paging VIOS partition provides access to all of the paging space devices for the shared memory partitions. The paging space devices can be located in physical storage in the server or on a storage area network (SAN). When you assign two paging VIOS partitions to the shared memory pool, you can configure each paging VIOS partition to access paging space devices in one of the following ways:

- You can configure each paging VIOS partition to access independent paging space devices. Paging space devices that are accessed by only one paging VIOS partition, or independent paging space devices, can be located in physical storage in the server or on a SAN.
- You can configure both paging VIOS partitions to access the same, or common, paging space devices. In this configuration, the paging VIOS partitions provide redundant access to paging space devices. When one paging VIOS partition becomes unavailable, the hypervisor sends a request to the other paging VIOS partition to retrieve the data on the paging space device. Common paging space devices must be located on a SAN to enable symmetrical access from both paging VIOS partitions.
- You can configure each paging VIOS partition to access some independent paging space devices and some common paging space devices.

If you configure the shared memory pool with two paging VIOS partitions, you can configure a shared memory partition to use either a single paging VIOS partition or redundant paging VIOS partitions. When you configure a shared memory partition to use redundant paging VIOS partitions, you assign a primary paging VIOS partition and a secondary paging VIOS partition to the shared memory partition. The hypervisor uses the primary paging VIOS partition to access the shared memory partition's paging space device. At this point, the primary paging VIOS partition is the current paging VIOS partition for the shared memory partition. The current paging VIOS partition is the paging VIOS partition that the hypervisor uses at any point in time to access data in the paging space device that is assigned to the shared memory partition. If the primary paging VIOS partition becomes unavailable, the hypervisor uses the secondary paging VIOS partition to access the shared memory partition's paging space device. At this point, the secondary paging VIOS partition becomes the current paging VIOS partition for the shared memory partition and continues as the current paging VIOS partition even after the primary paging VIOS partition becomes available again.

You do not need to assign the same primary and secondary paging VIOS partitions to all of the shared memory partitions. For example, you assign paging VIOS partition A and paging VIOS partition B to the shared memory pool. For one shared memory partition, you can assign paging VIOS partition A as the primary paging VIOS partition and paging VIOS partition B as the secondary paging VIOS partition. For a different shared memory partition, you can assign paging VIOS partition B as the primary paging VIOS partition and paging VIOS partition A as the secondary paging VIOS partition.

The following figure shows an example of a system with four shared memory partitions, two paging VIOS partitions, and four paging space devices.



The example shows the configuration options for paging VIOS partitions and paging space devices as described in the following table.

Table 5. Examples of paging VIOS partition configurations

Configuration option	Example
The paging space device that is assigned to a shared memory partition is located in physical storage in the server and is accessed by a single paging VIOS partition.	Paging space device 4 provides the paging space for Shared memory partition 4. Shared memory partition 4 is assigned to use Paging VIOS partition 2 to access Paging space device 4. Paging space device 4 is located in physical storage in the server and is assigned to Paging VIOS partition 2. Paging VIOS partition 2 is the only paging VIOS partition that can access Paging space device 4 (This relationship is shown by the blue line that connects Paging VIOS partition 2 to Paging space device 4.).
The paging space device that is assigned to a shared memory partition is located on a SAN and is accessed by a single paging VIOS partition.	Paging space device 1 provides the paging space for Shared memory partition 1. Shared memory partition 1 is assigned to use Paging VIOS partition 1 to access Paging space device 1. Paging space device 1 is connected to the SAN. Paging VIOS partition 1 is also connected to the SAN and is the only paging VIOS partition that can access Paging space device 1 (This relationship is shown by the green line that connects Paging VIOS partition 1 to Paging space device 1.).

Table 5. Examples of paging VIOS partition configurations (continued)

Configuration option	Example
The paging space device that is assigned to a shared memory partition is located on a SAN and is accessed redundantly by two paging VIOS partitions.	Paging space device 2 provides the paging space for Shared memory partition 2. Paging space device 2 is connected to the SAN. Paging VIOS partition 1 and Paging VIOS partition 2 are also connected to the SAN and can both access Paging space device 2. (These relationships are shown by the green line that connects Paging VIOS partition 1 to Paging space device 2 and the blue line that connects Paging VIOS partition 2 to Paging space device 2.) Shared memory partition 2 is assigned to use redundant paging VIOS partitions to access Paging space device 2. Paging VIOS partition 1 is configured as the primary paging VIOS partition and Paging VIOS partition 2 is configured as the secondary paging VIOS partition. Similarly, Paging space device 3 provides the paging space for Shared memory partition 3. Paging space device 3 is connected to the SAN. Paging VIOS partition 1 and Paging VIOS partition 2 are also connected to the SAN and can both access Paging space device 3. (These relationships are shown by the green line that connects Paging VIOS partition 1 to Paging space device 3 and the blue line that connects Paging VIOS partition 2 to Paging space device 3.) Shared memory partition 3 is assigned to use redundant paging VIOS partitions to access Paging space device 3. Paging VIOS partition 2 is configured as the primary paging VIOS partition and Paging VIOS partition 1 is configured as the secondary paging VIOS partition. Because Paging VIOS partition 1 and Paging VIOS partition 2 both have access to Paging space device 2 and Paging space device 3, Paging space device 2 and Paging space device 3 are common paging space devices that are accessed redundantly by Paging VIOS partition 1 and Paging VIOS partition 2. If Paging VIOS partition 1 becomes unavailable and Shared memory partition 2 needs to access data on its paging space device, the hypervisor sends a request to Paging VIOS partition 2 to retrieve the data on Paging space device 2. Similarly, if Paging VIOS partition 2 becomes unavailable and Shared memory partition 3 needs to access the data on its paging space device, the hypervisor sends a request to Paging VIOS partition 1 to retrieve the data on Paging space device 3.

Table 5. Examples of paging VIOS partition configurations (continued)

Configuration option	Example
A paging VIOS partition accesses both independent and common paging space devices.	Paging space device 1 and Paging space device 4 are independent paging space devices because only one paging VIOS partition accesses each. Paging VIOS partition 1 accesses Paging space device 1, and Paging VIOS partition 2 accesses Paging space device 4. Paging space device 2 and paging space device 3 are common paging space devices because both paging VIOS partitions access each. (These relationships are shown by the green and blue lines that connect the paging VIOS partitions to the paging space devices.) Paging VIOS partition 1 accesses the independent paging space device Paging space device 1, and also accesses the common paging space devices Paging space device 2 and Paging space device 3. Paging VIOS partition 2 accesses the independent paging space device Paging space device 4 and also accesses the common paging space devices Paging space device 2 and Paging space device 3.

When a single paging VIOS partition is assigned to the shared memory pool, you must shut down the shared memory partitions before you shut down the paging VIOS partition so that the shared memory partitions are not suspended when they attempt to access their paging space devices. When two paging VIOS partitions are assigned to the shared memory pool and the shared memory partitions are configured to use redundant paging VIOS partitions, you do not need to shut down the shared memory partitions to shut down a paging VIOS partition. When one paging VIOS partition is shut down, the shared memory partitions use the other paging VIOS partition to access their paging space devices. For example, you can shut down a paging VIOS partition and install VIOS updates without shutting down the shared memory partitions.

You can configure multiple VIOS logical partitions to provide access to paging space devices. However, you can only assign up to two of those VIOS partitions to the shared memory pool at any given time.

After you configure the shared memory partitions, you can later change the redundancy configuration of the paging VIOS partitions for a shared memory partition by modifying the partition profile of the shared memory partition and restarting the shared memory partition with the modified partition profile:

- You can change which paging VIOS partitions are assigned to a shared memory partition as the primary and secondary paging VIOS partitions.
- You can change the number of paging VIOS partitions that are assigned to a shared memory partition.

Paging space device

You can learn about how the Hardware Management Console (HMC) allocates and manipulates paging space devices on systems that use shared memory.

A *paging space device* is a physical or logical device that is used by a Virtual I/O Server to provide the paging space for a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*). The *paging space* is an area of nonvolatile storage used to hold portions of the shared memory partition's memory that are not resident in the shared memory pool.

Related concepts

Data flow for shared memory partitions

When the operating system that runs in a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) needs to access data, the data must reside in the shared memory pool. Systems with overcommitted memory configurations require the hypervisor and at least one Virtual I/O

Server (VIOS) logical partition that is assigned to the shared memory pool (hereafter referred to as *paging VIOS partition*) to move data between the shared memory pool and the paging space devices as needed.

Logical memory

Logical memory is the address space, assigned to a logical partition, that the operating system perceives as its main storage. For a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*), a subset of the logical memory is backed up by physical main storage and the remaining logical memory is kept in auxiliary storage.

Paging space devices on systems that are managed by an HMC

Learn about the location requirements, size requirements, and redundancy preferences for paging space devices on systems that are managed by a Hardware Management Console (HMC).

When you configure the shared memory pool, you assign can paging space devices to the shared memory pool. Paging space devices can be located in physical storage in the server or on a storage area network (SAN) as follows:

- Paging space devices that are accessed by a single Virtual I/O Server (VIOS) logical partition (hereafter referred to as a *paging VIOS partition*) can be located in physical storage in the server or on a SAN.
- Paging space devices that are accessed redundantly by two paging VIOS partitions, or *common* paging space devices, must be located on a SAN.

When you activate a shared memory partition, the HMC allocates a paging space device (that is assigned to the shared memory pool) to the shared memory partition. The HMC allocates only one paging space device to a shared memory partition at a time. When you shut down a shared memory partition, its paging space device becomes available to the HMC to allocate elsewhere. Thus, the fewest number of paging space devices that must be assigned to the shared memory pool is equal to the number of shared memory partitions that you plan to run simultaneously. After you create the shared memory pool, you can add or remove paging space devices to or from the shared memory pool as needed.

The HMC assigns paging space devices to shared memory partitions based on the size requirements for the shared memory partition and the redundancy preferences that you specify for partition activation.

Size requirements

The HMC allocates a paging space device to a shared memory partition that best fits the size requirements of the shared memory partition.

- For AIX and Linux shared memory partitions, the paging space device must be at least the size of the maximum logical memory size of the shared memory partition.
- For IBM i shared memory partitions, the paging space device must be at least the size of the maximum logical memory size of the shared memory partition plus 8 KB for every megabyte.

Shared memory partitions might have several partition profiles that specify different maximum logical memory sizes. To maintain flexibility, consider creating paging space devices that are large enough to be used by shared memory partitions with multiple partition profiles. When you activate a shared memory partition with a different partition profile, the shared memory partition already has a paging space device allocated to it based on the size requirements of the previously activated partition profile. If you create a paging space device that is large enough to meet the size requirements of multiple partition profiles and you activate the shared memory partition with a different partition profile, the HMC can use the same paging space device for the newly activated partition profile. If the paging space device does not meet the size requirements of the newly activated partition profile, the HMC frees the paging space device currently allocated to the shared memory partition and allocates a different paging space device that meets the size requirements specified in the newly activated partition profile.

Redundancy preferences

The HMC allocates a paging space device to a shared memory partition that satisfies the redundancy preferences that you specify for partition activation:

- If you specify that the shared memory partition uses redundant paging VIOS partitions, the HMC uses the following process to select a suitable paging space device for the shared memory partition:
 1. The HMC assigns a paging space device that is common and available. (A paging space device is *available* when it is not currently assigned to a shared memory partition and is inactive.)
 2. If the HMC cannot find a paging space device that is common and available, it reassigns a paging space device that is common and unavailable. (A paging space device is *unavailable* when it is active and currently assigned to a shared memory partition that is shut down.)
 3. If the HMC cannot find a paging space device that is common and unavailable, it cannot activate the shared memory partition.
- If you specify that the shared memory partition does not use redundant paging VIOS partitions, the HMC uses the following process to select a suitable paging space device for the shared memory partition:
 1. The HMC assigns a paging space device that is independent and available. (A paging space device is *independent* when it is accessed by only one paging VIOS partition that is assigned to the shared memory partition.)
 2. If the HMC cannot find a paging space device that is independent and available, the HMC reassigns a paging space device that is independent and unavailable.
 3. If the HMC cannot find a paging space device that is independent and unavailable, and two paging VIOS partitions are assigned to the shared memory pool, then the HMC assigns a paging space device that is common and available. In this situation, the shared memory partition does not use redundant paging VIOS partitions even though its paging space device can be accessed by both paging VIOS partitions. Also, the partition profile need not specify the second paging VIOS partition.
 4. If the HMC cannot find a paging space device that is common and available, and two paging VIOS partitions are assigned to the shared memory pool, then the HMC reassigns a paging space device that is common and unavailable. In this situation, the shared memory partition does not use redundant paging VIOS partitions even though its paging space device can be accessed by both paging VIOS partitions. Also, the partition profile need not specify the second paging VIOS partition.
 5. If the HMC cannot find a paging space device that is common and unavailable, it cannot activate the shared memory partition.
- If you specify that the shared memory partition use redundant paging VIOS partitions, if possible, the HMC uses the following process to select a suitable paging space device for the shared memory partition:
 1. The HMC assigns a paging space device that is common and available.
 2. If the HMC cannot find a paging space device that is common and available, it assigns a paging space device that is common and unavailable.
 3. If the HMC cannot find a paging space device that is common and unavailable, it assigns a paging space device that is independent and available to the primary paging VIOS partition. In this situation, the shared memory partition does not use redundant paging VIOS partitions and the primary paging VIOS partition is the only paging VIOS partition that is assigned to the shared memory partition.
 4. If the HMC cannot find a paging space device that is independent and available to the primary paging VIOS partition, it assigns a paging space device that is independent and unavailable to the primary paging VIOS partition. In this situation, the shared memory partition does not use redundant paging VIOS partitions and the primary paging VIOS partition is the only paging VIOS partition that is assigned to the shared memory partition.
 5. If the HMC cannot find a paging space device that is independent and unavailable to the primary paging VIOS partition, it assigns a paging space device that is independent and available to the secondary paging VIOS partition. In this situation, the shared memory partition does not use redundant paging VIOS partitions and the secondary paging VIOS partition is the only paging VIOS partition that is assigned to the shared memory partition.
 6. If the HMC cannot find a paging space device that is independent and available to the secondary paging VIOS partition, it assigns a paging space device that is independent and unavailable to the secondary paging VIOS partition. In this situation, the shared memory partition does not use

redundant paging VIOS partitions and the secondary paging VIOS partition is the only paging VIOS partition that is assigned to the shared memory partition.

7. If the HMC cannot find a paging space device that is independent and unavailable to the secondary paging VIOS partition, it cannot activate the shared memory partition.

Related concepts

Partition profile

A partition profile is a record on the Hardware Management Console (HMC) that specifies a possible configuration for a logical partition. When you activate a logical partition by using a partition profile, the managed system attempts to start the logical partition by using the configuration information in the partition profile.

Related tasks

Preparing to configure shared memory

Before you configure the shared memory pool and create logical partitions that use shared memory (hereafter referred to as *shared memory partitions*), you need to plan for the shared memory pool, the shared memory partitions, the paging space devices, and the Virtual I/O Server logical partitions (hereafter referred to as *paging VIOS partitions*).

Related reference

Configuration requirements for shared memory

Review the requirements for the system, Virtual I/O Server (VIOS), logical partitions, and paging space devices so that you can successfully configure shared memory.

Shared memory distribution

The hypervisor uses the memory weight of each logical partition that uses shared memory (hereafter referred to as *shared memory partitions*) to help determine which logical partitions receive more physical memory from the shared memory pool. To help optimize performance and memory use, the operating systems that run in shared memory partitions provide the hypervisor with information about how the operating system uses its memory to help the hypervisor determine which pages to store in the shared memory pool and which pages to store in the paging space devices.

In a shared memory configuration that is physically overcommitted (where the sum of the logical memory that is currently used by all shared memory partitions is greater than the amount of memory in the shared memory pool), the hypervisor stores a portion of the logical memory in the shared memory pool and stores the remainder of the logical memory in the paging space devices. The hypervisor determines the amount of physical memory to allocate from the shared memory pool to each shared memory partition and the amount of logical memory to store in the paging space devices. The hypervisor also determines which pieces, or pages, of memory to store in each location.

The smallest amount of physical memory that the hypervisor can allocate from the shared memory pool to a shared memory partition at any given time is the amount of physical memory that the shared memory partition requires for its I/O devices. The hypervisor guarantees to each shared memory partition that the shared memory partition can use a portion of the shared memory pool for its I/O devices, up to the I/O entitled memory that is assigned to the shared memory partition. The largest amount of physical memory that the hypervisor can allocate from the shared memory pool to a shared memory partition at any given time is the amount of logical memory assigned to the shared memory partition.

The amount of physical memory from the shared memory pool that the hypervisor allocates to the shared memory partitions is determined by the workloads that are running in the shared memory partitions and the amount of logical memory that is assigned to each shared memory partition. You can influence how much physical memory the hypervisor allocates from the shared memory pool to each shared memory partition by specifying a memory weight for each shared memory partition. *Memory weight* is a relative value that is one of the factors that the hypervisor uses to allocate physical memory from the shared memory pool to the shared memory partitions. A higher memory weight relative to the memory weights of other shared memory partitions increases the probability that the hypervisor allocates more physical memory to a shared memory partition.

To help maintain the best possible performance, the operating system that runs in a shared memory partition continually attempts to operate within the amount of physical memory allocated to it from the shared memory pool by moving its overcommitted logical memory to a paging space. In general, the

operating system moves its memory to a paging space more often when it runs in a shared memory partition than when it runs in a dedicated memory partition. Therefore, the paging space that the operating system uses to manage its memory needs to be larger when the logical partition uses shared memory than when the logical partition uses dedicated memory.

The operating systems that run in shared memory partitions provide information to the hypervisor about how the operating system uses its pages. When the hypervisor manages the overcommitted logical memory, it uses this information to determine which pages to store in the paging space device and which pages to store in the shared memory pool. When the hypervisor needs to deallocate physical memory from the shared memory partition and move it to the paging space device, the hypervisor requests the operating system to release pages. The operating system might mark the pages that it will not use, and the hypervisor moves the marked pages first. This enables the hypervisor to select the most optimal pages to move out of the shared memory pool, which improves memory use and performance. For example, the operating system uses one page for kernel data and another page for cache and the hypervisor needs to move one page to the paging space device. The hypervisor moves the cache page to the paging space device to optimize performance.

Related concepts

Paging space device

You can learn about how the Hardware Management Console (HMC) allocates and manipulates paging space devices on systems that use shared memory.

Data flow for shared memory partitions

When the operating system that runs in a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) needs to access data, the data must reside in the shared memory pool. Systems with overcommitted memory configurations require the hypervisor and at least one Virtual I/O Server (VIOS) logical partition that is assigned to the shared memory pool (hereafter referred to as *paging VIOS partition*) to move data between the shared memory pool and the paging space devices as needed.

Performance considerations for over committed shared memory partitions

Learn about how the degree to which the memory configuration of a logical partition that uses shared memory (also referred to as a *shared memory partition*) is over committed affects the performance of the shared memory partition. In general, the less over committed the memory configuration of a shared memory partition, the better its performance.

Logical memory

Logical memory is the address space, assigned to a logical partition, that the operating system perceives as its main storage. For a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*), a subset of the logical memory is backed up by physical main storage and the remaining logical memory is kept in auxiliary storage.

Active Memory Expansion for AIX logical partitions

When you enable Active Memory Expansion for an AIX logical partition, you increase the memory capacity of the logical partition without assigning more memory to it. The operating system compresses a portion of the memory that the logical partition uses. This compression creates space for more data and expanding the memory capacity of the logical partition.

When you expand the memory capacity of a logical partition, you enable the logical partition to do more work with the same amount of memory. This can be especially useful when you want to increase the workload of a logical partition, but cannot assign more memory to the logical partition. When you expand the memory capacity of several logical partitions on a server, you increase the overall memory capacity of the server. This can be especially useful when you want to consolidate more workloads onto the server by creating more logical partitions.

You can configure the degree of memory expansion that you want to achieve for a logical partition by setting the Active Memory Expansion factor in a partition profile of the logical partition. The expansion factor is a multiplier of the amount of memory that is assigned to the logical partition. For example, if the amount of memory that is assigned to a logical partition is 25 GB and the expansion factor is set to 2.0, then the desired memory capacity of the logical partition is 50 GB. In this situation, the operating system attempts to compress data such that 50 GB of data fits into 25 GB of memory. After you set the

expansion factor, you can monitor the performance of the logical partition and then dynamically change the expansion factor to improve performance.

When you configure a logical partition to use Active Memory Expansion, you must also configure some additional processing resources for the logical partition. The operating system uses the additional processing resources to perform the memory compression. The amount of processing resources that the logical partition requires depends on the workload that is running in the logical partition and the expansion factor that you set for the logical partition.

You can configure Active Memory Expansion for logical partitions that use dedicated memory and logical partitions that use shared memory.

Related information

[IBM AIX Knowledge Center website](#)

Terminal and console options for logical partitions

You can initiate a terminal or console session to the logical partitions on your managed system using various methods. Your choice of terminal or console depends on your operating system and business needs.

The following choices of terminal or console are available for each operating system.

<i>Table 6. Terminal and console options for logical partitions</i>	
Operating system	Terminal or console options
AIX	• Hardware Management Console (HMC) • Telnet • OpenSSH with OpenSSL (included in the AIX expansion pack) • Direct serial connection (ASCII terminal or PC connected with null modem cable) • IBM i virtual console (for AIX logical partitions that use IBM i resources) • When on a system with a Virtual I/O Server (VIOS) logical partition, the console can be provided by the VIOS logical partition when using VIOS 1.2.0 or later.
Linux	• HMC • Telnet • OpenSSH with OpenSSL (included in the Linux distribution) • Direct serial connection (ASCII terminal or PC connected with null modem cable) • IBM i virtual console (for Linux logical partitions that use IBM i resources) • When on a system with a Virtual I/O Server (VIOS) logical partition, the console can be provided by the VIOS logical partition when using VIOS 1.2.0 or later.
Virtual I/O Server	• Hardware Management Console (HMC) • Telnet • OpenSSH with OpenSSL (included in the AIX expansion pack) • Direct serial connection (ASCII terminal or PC connected with null modem cable) • IBM i virtual console (for AIX logical partitions that use IBM i resources) • When on a system with a Virtual I/O Server (VIOS) logical partition, the console can be provided by the VIOS logical partition when using VIOS 1.2.0 or later.


Hardware Management Console terminal and console options

The HMC provides virtual terminal emulation for AIX and Linux logical partitions and virtual 5250 console emulation for IBM i logical partitions.

You can create virtual terminal and virtual 5250 console sessions locally on the HMC by using the Server Management commands on the HMC. If you configure the HMC to allow remote access, you can also create virtual terminal and virtual 5250 console sessions remotely through the HMC. You can create remote virtual terminal sessions on AIX and Linux logical partitions by using the Server Management commands. You can also create virtual 5250 console sessions on IBM i logical partitions. You must configure the HMC to allow remote access, and you must configure encryption on the logical partitions for the session to be secure.

The HMC communicates with servers by using service applications to detect, consolidate, and send information to IBM for analysis.

The following figure shows a partitioned server being managed by an HMC.



Operations Console for IBM i logical partitions


Operations Console allows you to use a local or remote PC to access IBM i on logical partitions. Operations Console is an installable component of IBM i Access for Windows licensed program.

You can use Operations Console to connect to IBM i logical partitions over a network (LAN connection).

The management tasks that you can perform using Operations Console depend upon whether you are managing your logical partitions using a Hardware Management Console (HMC) or using the Virtual Partition Manager on IBM i:

- If you are using an HMC to manage your logical partitions, you can use Operations Console to access IBM i on your IBM i logical partitions.
- If you are using the Virtual Partition Manager on IBM i to manage your logical partitions, you can use Operations Console to access IBM i on your IBM i logical partitions. In turn, you can use Operations Console to access the Virtual Partition Manager on your IBM i logical partitions. This allows you to create up to four Linux logical partitions on the managed system and manage the resources for all logical partitions on the managed system.

The following figure shows a partitioned server with an HMC and local console on a network.



Related information

[Managing Operations](#)

I/O devices

I/O devices allow your managed system to gather, store, and transmit data. I/O devices are found in the server unit itself and in expansion units that are attached to the server. I/O devices can be embedded into the unit, or they can be installed into physical slots.

Not all types of I/O devices are supported for all operating systems or on all server models. For example, Switch Network Interface (SNI) adapters are supported only on certain server models, and are not supported for IBM i logical partitions.

Single root I/O virtualization (SR-IOV) allows virtualization of the physical ports of an adapter so that the ports can be shared by multiple partitions that are running simultaneously. To share the ports of an SR-IOV capable adapter, the adapter must first be enabled for the SR-IOV shared mode. After an adapter is enabled for SR-IOV shared mode, SR-IOV logical ports can be assigned to logical partitions.



Attention: Some PCI adapters and embedded controllers require multiple PCI or PCI-E slots to be associated with them. Carefully review the PCI or PCI-E slot assignments for each logical partition to ensure that the slot configuration of the logical partition meets the adapter functional requirements. For details, see [Managing PCI adapters](#), and [PCI adapter placement](#).

Virtual adapters

With virtual adapters, you can connect logical partitions with each other without using physical hardware. Operating systems can display, configure, and use virtual adapters just like they can display, configure, and use physical adapters. Depending on the operating environment used by the logical partition, you can create virtual Ethernet adapters, virtual Fibre Channel adapters, virtual Small Computer Serial Interface (SCSI) adapters, and virtual serial adapters for a logical partition.

The system administrator uses the following tools to create virtual adapters:

- Hardware Management Console (HMC)
- Virtual Partition Manager

Adapters can be added while the system is running using dynamic partitioning. The virtual adapters are recorded in system inventory and management utilities. Converged location codes can be used to correlate operating-system level or partition-level software entities to adapters, such as eth0, CMN21, and en0. Similarly, the Ethernet adapters are visible in the same way as physical Ethernet adapters.

By default, virtual Ethernet Media Access Control (MAC) addresses are created from the locally administered range. Using the default MAC addresses, it is possible that different servers will have virtual Ethernet adapters with the same addresses. This situation can present a problem if multiple, virtual networks are bridged to the same physical network.

If a server logical partition providing I/O for a client logical partition fails, the client logical partition might continue to function, depending on the significance of the hardware it is using. For example, if one logical partition is providing the paging volume for another logical partition, a failure of the logical partition providing that particular resource will be significant to the other logical partition. However, if the shared resource is a tape drive, a failure of the server logical partition providing the resource will have only minimal effects on the client logical partition.

Client support for virtual I/O

The following table summarizes operating system support for using virtual I/O devices.

Table 7. Client support for virtual I/O by operating system						
Client operating system	Virtual console	Virtual Ethernet	Virtual Fibre Channel	Virtual disk	Virtual optical	Virtual tape
AIX	Yes	Yes	Yes	Yes	Yes, on HMC-managed systems, at least one Virtual I/O Server logical partition must be present	Yes
AIX	Yes	Yes	Yes	Yes	Yes, on HMC-managed systems, at least one Virtual I/O Server logical partition must be present	Yes
IBM i	Yes	Yes	Yes	Yes	Yes	Yes
Linux	Yes	Yes	Yes	Yes	Yes	Yes

Table 7. Client support for virtual I/O by operating system (continued)

Client operating system	Virtual console	Virtual Ethernet	Virtual Fibre Channel	Virtual disk	Virtual optical	Virtual tape
Linux	Yes	Yes	Yes	Yes	Yes	Yes

AIX logical partitions support booting from virtual devices, including disk boot from virtual disk, network boot from virtual Ethernet, and tape boot from virtual tape.

The firmware running in AIX and Linux logical partitions recognizes virtual I/O and can start the logical partition from virtual I/O. An IPL can be performed either from the network over virtual Ethernet or from a device such as virtual disk or virtual CD.

Server support for virtual I/O

The following table summarizes operating system support for providing virtual I/O to logical partitions.

Table 8. Server support for virtual I/O by operating system

Server	Virtual optical	Virtual console	Virtual disk	Virtual tape	Virtual Fibre Channel
IBM i	Yes	Yes	Yes	Yes	No
Linux	Yes	Yes	No	No	No
Linux	Yes	Yes	No	No	No
Virtual I/O Server	Yes	Yes	Yes	Yes	Yes

Virtual I/O Server provides SCSI disk, shared Ethernet, virtual Fibre Channel, virtual optical, and virtual tape functions to logical partitions that use Virtual I/O Server resources. On VIOS Version 2.2.0.11, Fix Pack 24, Service Pack 1, or later, you can create a cluster of only one Virtual I/O Server (VIOS) partition connected to the same shared storage pool and that has access to distributed storage. On VIOS Version 2.2.1.3, or later, you can create a cluster that consists of up to four VIOS partitions. The Virtual I/O Server also provides a virtual console to AIX and Linux logical partitions.

IBM i provides disk, CD, tape, and console functions to logical partitions that use IBM i resources. IBM i uses standard network server storage and network server descriptions to provide disk, CD, and tape resources to other logical partitions. An IBM i logical partition cannot simultaneously provide virtual resources to other logical partitions and use virtual resources provided by another IBM i logical partition or by the Virtual I/O Server logical partition.

To configure virtual I/O for the logical partitions on your managed system, you must create virtual I/O adapters on the HMC. Virtual I/O adapters are usually created when you create your logical partitions. Alternatively, you can add virtual I/O adapters to running logical partitions using dynamic partitioning. After you create a virtual I/O adapter, you can then access the operating system used by the logical partition and complete the configuration of the virtual I/O adapter in the operating system software. For Linux partitions, virtual adapters are listed in the device tree. The device tree contains virtual SCSI adapters, not the devices under the adapter.

Logical Host Ethernet Adapter

A logical Host Ethernet Adapter (LHEA) is a special type of virtual adapter. Even though an LHEA is a virtual resource, an LHEA can exist only if a physical Host Ethernet Adapter, or Integrated Virtual Ethernet, provides its resources to the LHEA.

Note: HEA is not supported on POWER9 processor-based server.

Related concepts

I/O device assignment in partition profiles

I/O devices are assigned to partition profiles either on a slot-by-slot basis, or on logical port basis in the case of shared mode single root I/O virtualization (SR-IOV) adapters. For I/O devices that are assigned to partition profiles on a slot-by-slot basis, most I/O devices can be assigned to a partition profile on the HMC as required or as allocated. For SR-IOV logical ports, I/O devices are always assigned to a profile as required.

Host Ethernet Adapter

A *Host Ethernet Adapter (HEA)* is a physical Ethernet adapter that is integrated directly into the GX+ bus on a managed system. HEAs offer high throughput, low latency, and virtualization support for Ethernet connections. HEAs are also known as Integrated Virtual Ethernet adapters (IVE adapters).

Virtual Ethernet

Virtual Ethernet allows logical partitions to communicate with each other without having to assign physical hardware to the logical partitions.

You can create virtual Ethernet adapters on each logical partition and connect these virtual Ethernet adapters to virtual LANs. TCP/IP communications over these virtual LANs is routed through the server firmware.

A virtual Ethernet adapter provides similar function as a 1 Gb Ethernet adapter. A logical partition can use virtual Ethernet adapters to establish multiple high-speed interpartition connections within a single managed system. AIX, IBM i, Linux, and Virtual I/O Server logical partitions can communicate with each other using TCP/IP over the virtual Ethernet communications ports.

Virtual Ethernet adapters are connected to an IEEE 802.1q (VLAN)-style virtual Ethernet switch. Using this switch function, logical partitions can communicate with each other by using virtual Ethernet adapters and assigning VLAN IDs that enable them to share a common logical network. The virtual Ethernet adapters are created and the VLAN ID assignments are done using the Hardware Management Console (HMC). The system transmits packets by copying the packet directly from the memory of the sender logical partition to the receive buffers of the receiver logical partition without any intermediate buffering of the packet.

You can configure an Ethernet bridge between the virtual LAN and a physical Ethernet adapter that is owned by a Virtual I/O Server or IBM i logical partition. The logical partitions on the virtual LAN can communicate with an external Ethernet network through the Ethernet bridge. You can reduce the number of physical Ethernet adapters required for a managed system by routing external communications through the Ethernet bridge.

The number of virtual Ethernet adapters allowed for each logical partition varies by operating system.

- AIX 5.3 and later supports up to 256 virtual Ethernet adapters for each logical partition.
- Version 2.6 of the Linux kernel supports up to 32, 768 virtual Ethernet adapters for each logical partition. Each Linux logical partition can belong to a maximum of 4, 094 virtual LANs.

Besides a Port VLAN ID, the number of additional VLAN ID values that can be assigned for each virtual Ethernet adapter is 19, which indicates that each virtual Ethernet adapter can be used to access 20 networks. The HMC generates a locally administered Ethernet MAC address for the virtual Ethernet adapters so that these addresses do not conflict with physical Ethernet adapter MAC addresses.

After a specific virtual Ethernet is enabled for a logical partition, a network device is created in the logical partition. This network device is named `entX` on AIX logical partitions and `ethX` on Linux logical partitions, where `X` represents sequentially assigned numbers. The user can then set up TCP/IP configuration similar to a physical Ethernet device to communicate with other logical partitions.

If a virtual Ethernet adapter is set for checksum offload, the virtual Ethernet adapter cannot generate a checksum for any packet that the virtual Ethernet adapter sends to a multicast or broadcast MAC address.

Some managed systems contain a Host Ethernet Adapter (HEA). A *Host Ethernet Adapter (HEA)* is a physical Ethernet adapter that is integrated directly into the GX+ bus on a managed system. HEAs are also known as Integrated Virtual Ethernet adapters (IVE adapters). Unlike most other types of I/O devices, you can never assign the HEA itself to a logical partition. Instead, multiple logical partitions can connect

directly to the HEA and use the HEA resources. This allows these logical partitions to access external networks through the HEA using an Ethernet bridge on another logical partition.

Note: HEA is not supported on POWER9 processor-based server.

You can enable and disable individual virtual Ethernet adapters by using the Hardware Management Console (HMC). You can use the **chhwres** command to enable or disable a virtual Ethernet adapter. A particular logical partition can be removed from the network when the virtual Ethernet adapter is disabled. You can reconnect the logical partition to the network by enabling the virtual Ethernet adapter. To reconnect the logical partition, you must use a virtual Ethernet that is bridged by using a Shared Ethernet Adapter (SEA) in the Virtual I/O Server (VIOS). The status of the virtual Ethernet adapter can be queried at any time by using the **lshwres** command. The disabled state persists during partition restart. Trunk adapters cannot be disabled. You must have super administrator or product engineer access to the HMC to enable or disable a virtual Ethernet adapter.

Related concepts

Host Ethernet Adapter

A *Host Ethernet Adapter (HEA)* is a physical Ethernet adapter that is integrated directly into the GX+ bus on a managed system. HEAs offer high throughput, low latency, and virtualization support for Ethernet connections. HEAs are also known as Integrated Virtual Ethernet adapters (IVE adapters).

Virtual Fibre Channel


With N_Port ID Virtualization (NPIV), you can configure the managed system so that multiple logical partitions can access independent physical storage through the same physical Fibre Channel adapter.

To access physical storage in a typical storage area network (SAN) that uses Fibre Channel, the physical storage is mapped to logical units (LUNs) and the LUNs are mapped to the ports of physical Fibre Channel adapters. Each physical port on each physical Fibre Channel adapter is identified using one worldwide port name (WWPN).

NPIV is a standard technology for Fibre Channel networks that enables you to connect multiple logical partitions to one physical port of a physical Fibre Channel adapter. Each logical partition is identified by a unique WWPN, which means that you can connect each logical partition to independent physical storage on a SAN.

To enable NPIV on the managed system, you must create a Virtual I/O Server logical partition (version 2.1, or later) that provides virtual resources to client logical partitions. You assign the physical Fibre Channel adapters (that support NPIV) to the Virtual I/O Server logical partition. Then, you can connect virtual Fibre Channel adapters on the client logical partitions to virtual Fibre Channel adapters on the Virtual I/O Server logical partition. A *virtual Fibre Channel adapter* is a virtual adapter that provides client logical partitions with a Fibre Channel connection to a storage area network through the Virtual I/O Server logical partition. The Virtual I/O Server logical partition provides the connection between the virtual Fibre Channel adapters on the Virtual I/O Server logical partition and the physical Fibre Channel adapters on the managed system.

The following figure shows a managed system configured to use NPIV.



P9HAT008-0

The figure shows the following connections:

- A storage area network (SAN) connects three units of physical storage to a physical Fibre Channel adapter that is located on the managed system. The physical Fibre Channel adapter is assigned to the Virtual I/O Server and supports NPIV.
- The physical Fibre Channel adapter connects to three virtual Fibre Channel adapters on the Virtual I/O Server. All three virtual Fibre Channel adapters on the Virtual I/O Server connect to the same physical port on the physical Fibre Channel adapter.
- Each virtual Fibre Channel adapter on the Virtual I/O Server connects to one virtual Fibre Channel adapter on a client logical partition. Each virtual Fibre Channel adapter on each client logical partition receives a pair of unique WWPNs. The client logical partition uses one WWPN to log into the SAN at any given time. The other WWPN is used when you move the client logical partition to another managed system.

Using their unique WWPNs and the virtual Fibre Channel connections to the physical Fibre Channel adapter, the operating systems that run in the client logical partitions discover, instantiate, and manage their physical storage located on the SAN. In the previous figure, Client logical partition 1 accesses Physical storage 1, Client logical partition 2 accesses Physical storage 2, and Client logical partition 3 accesses Physical storage 3. For IBM i client partitions, the LUNs of the physical storage connected with NPIV require a storage-specific device driver and do not use the generic virtual SCSI device driver. The Virtual I/O Server cannot access and does not emulate the physical storage to which the client logical partitions have access. The Virtual I/O Server provides the client logical partitions with a connection to the physical Fibre Channel adapters on the managed system.

There is always a one-to-one relationship between virtual Fibre Channel adapters on the client logical partitions and the virtual Fibre Channel adapters on the Virtual I/O Server logical partition. That is, each

virtual Fibre Channel adapter on a client logical partition must connect to only one virtual Fibre Channel adapter on the Virtual I/O Server logical partition, and each virtual Fibre Channel on the Virtual I/O Server logical partition must connect to only one virtual Fibre Channel adapter on a client logical partition. Mapping of multiple Virtual Fibre Channel adapters of a single client logical partition through multiple virtual server Fibre Channel adapters to the same physical Fibre Channel adapter is not recommended.

Using SAN tools, you can zone and mask LUNs that include WWPNs that are assigned to virtual Fibre Channel adapters on client logical partitions. The SAN uses WWPNs that are assigned to virtual Fibre Channel adapters on client logical partitions the same way it uses WWPNs that are assigned to physical ports.

Related information

[Redundancy configuration using Virtual Fibre channel adapters](#)

Virtual Fibre Channel for HMC-managed systems

On systems that are managed by the Hardware Management Console (HMC), you can dynamically add and remove virtual Fibre Channel adapters to and from the Virtual I/O Server logical partition and each client logical partition. You can also view information about the virtual and physical Fibre Channel adapters and the worldwide port names (WWPNs) by using Virtual I/O Server commands.

To enable N_Port ID Virtualization (NPIV) on the managed system, you create the required virtual Fibre Channel adapters and connections as follows:

- You use the HMC to create virtual Fibre Channel adapters on the Virtual I/O Server logical partition and associate them with virtual Fibre Channel adapters on the client logical partitions.
- You use the HMC to create virtual Fibre Channel adapters on each client logical partition and associate them with virtual Fibre Channel adapters on the Virtual I/O Server logical partition. When you create a virtual Fibre Channel adapter on a client logical partition, the HMC generates a pair of unique WWPNs for the client virtual Fibre Channel adapter.
- You can connect the virtual Fibre Channel adapters on the Virtual I/O Server to the physical ports of the physical Fibre Channel adapter by running the **vfcmap** command on the Virtual I/O Server.

The HMC generates WWPNs based on the range of names available for use with the prefix in the vital product data on the managed system. This 6–digit prefix comes with the purchase of the managed system and includes 32,000 pairs of WWPNs. When you remove a virtual Fibre Channel adapter from a client logical partition, the hypervisor deletes the WWPNs that are assigned to the virtual Fibre Channel adapter on the client logical partition. The HMC does not reuse the deleted WWPNs when generating WWPNs for virtual Fibre Channel adapters in the future. If you run out of WWPNs, you must obtain an activation code that includes another prefix with another 32,000 pairs of WWPNs.

To avoid configuring the physical Fibre Channel adapter to be a single point of failure for the connection between the client logical partition and its physical storage on the SAN, do not connect two virtual Fibre Channel adapters from the same client logical partition to the same physical Fibre Channel adapter. Instead, connect each virtual Fibre Channel adapter to a different physical Fibre Channel adapter.

You can dynamically add and remove virtual Fibre Channel adapters to and from the Virtual I/O Server logical partition and to and from client logical partitions.

<i>Table 9. Dynamic partitioning tasks and results for virtual Fibre Channel adapters</i>		
Dynamically add or remove virtual Fibre Channel adapter	To or from a client logical partition or a Virtual I/O Server logical partition	Result
Add a virtual Fibre Channel adapter	To a client logical partition	The HMC generates the a pair of unique WWPNs for the client virtual Fibre Channel adapter.
Add a virtual Fibre Channel adapter	To a Virtual I/O Server logical partition	You need to connect the virtual Fibre Channel adapter to a physical port on a physical Fibre Channel adapter.

Table 9. Dynamic partitioning tasks and results for virtual Fibre Channel adapters (continued)

Dynamically add or remove virtual Fibre Channel adapter	To or from a client logical partition or a Virtual I/O Server logical partition	Result
Remove a virtual Fibre Channel adapter	From a client logical partition	- The hypervisor deletes the WWPNs and does not reuse them. - You must either remove the associated virtual Fibre Channel adapter from the Virtual I/O Server, or associate it with another virtual Fibre Channel adapter on a client logical partition.
Remove a virtual Fibre Channel adapter	From a Virtual I/O Server logical partition	- The Virtual I/O Server removes the connection to the physical port on the physical Fibre Channel adapter. - You must either remove the associated virtual Fibre Channel adapter from the client logical partition, or associate it with another virtual Fibre Channel adapter on the Virtual I/O Server logical partition.

The following table lists the Virtual I/O Server commands that you can run to view information about the Fibre Channel adapters.

Table 10. Virtual I/O Server commands that display information about Fibre Channel adapters

Virtual I/O Server command	Information displayed by command
lsmap	- Displays the virtual Fibre Channel adapters on the Virtual I/O Server that are connected to the physical Fibre Channel adapter - Displays attributes of the virtual Fibre Channel adapters on the client logical partitions that are associated with the virtual Fibre Channel adapters on the Virtual I/O Server that are connected to the physical Fibre Channel adapter
lsnports	Displays information about the physical ports on the physical Fibre Channel adapters that support NPIV, such as: - The name and location code of the physical port - The number of available physical ports - The total number of WWPNs that the physical port can support - Whether the switches, to which the physical Fibre Channel adapters are cabled, support NPIV

You can also run the **lshwres** command on the HMC to display the remaining number of WWPNs and to display the prefix that is currently used to generate the WWPNs.

Virtual SCSI adapters

Virtual SCSI (Small Computer Systems Interface) adapters provide one logical partition with the ability to use storage I/O (disk, CD, and tape) that is owned by another logical partition.

A virtual SCSI client adapter in one logical partition can communicate with a virtual SCSI server adapter in another logical partition. The virtual SCSI client adapter allows a logical partition to access a storage device being made available by the other logical partition. The logical partition owning the hardware is the *server logical partition*, and the logical partition that uses the virtualized hardware is the *client logical partition*. With this arrangement, the system can have many server logical partitions.

For example, logical partition A provides disk space to logical partitions B, C, and D. A logical partition can simultaneously use virtual I/O from more than one logical partition. Therefore, using the example, while logical partition A provides disk space to logical partitions B, C, and D, logical partitions A and B can use a tape drive connected to logical partition D. In this case, A is serving D for disk space, while D is serving A for the tape device.

Virtual SCSI allows you to simplify the backup and maintenance operations on your managed system. When you back up the data on the server logical partition, you also back up the data on each client logical partition.

Virtual SCSI server adapters can be created only in logical partitions of type IBM i and Virtual I/O Server.

The virtual SCSI client device driver is not capable of storage protection using Redundant Arrays of Independent Disks (RAID). While the Linux operating system allows software RAID protection of virtual disks, the recommended technique for protecting disk storage is to configure the virtual I/O storage server to perform the disk protection.

For HMC-managed systems, virtual SCSI adapters are created and assigned to logical partitions using partition profiles.

Related concepts

Partition profile

A partition profile is a record on the Hardware Management Console (HMC) that specifies a possible configuration for a logical partition. When you activate a logical partition by using a partition profile, the managed system attempts to start the logical partition by using the configuration information in the partition profile.

Virtual serial adapters

Virtual serial adapters provide a point-to-point connection from one logical partition to another, or from the Hardware Management Console (HMC) to each logical partition on the managed system. Virtual serial adapters are used primarily to establish terminal or console connections to logical partitions.

When you create a logical partition, the HMC automatically creates two virtual server serial adapters on the logical partition. These virtual server serial adapters allow you to establish a terminal or console connection to the logical partition through the HMC.

You can also create pairs of virtual serial adapters on logical partitions so that you can access and control one logical partition directly from another logical partition. For example, one logical partition uses the disk resources of another logical partition using virtual SCSI adapters. You can create a server serial adapter on the logical partition that uses the disk resources and a client serial adapter on the logical partition that owns the disk resources. This connection allows the logical partition that owns the disk resources to shut down the logical partition that uses the disk resources before you back up data on the logical partition that owns the disk resources.

On HMC-managed systems, virtual serial adapters are created and assigned to logical partitions using partition profiles.

Related concepts

Partition profile

A partition profile is a record on the Hardware Management Console (HMC) that specifies a possible configuration for a logical partition. When you activate a logical partition by using a partition profile, the managed system attempts to start the logical partition by using the configuration information in the partition profile.

Host Ethernet Adapter

A *Host Ethernet Adapter (HEA)* is a physical Ethernet adapter that is integrated directly into the GX+ bus on a managed system. HEAs offer high throughput, low latency, and virtualization support for Ethernet connections. HEAs are also known as Integrated Virtual Ethernet adapters (IVE adapters).

Note: HEA is not supported on POWER9 processor-based server.

Unlike most other types of I/O devices, you can never assign the HEA itself to a logical partition. Instead, multiple logical partitions can connect directly to the HEA and use the HEA resources. This allows these logical partitions to access external networks through the HEA without having to go through an Ethernet bridge on another logical partition.

To connect a logical partition to an HEA, you must create a logical Host Ethernet Adapter (LHEA) for the logical partition. A *logical Host Ethernet Adapter (LHEA)* is a representation of a physical HEA on a logical partition. An LHEA appears to the operating system as if it were a physical Ethernet adapter, just as a virtual Ethernet adapter appears as if it were a physical Ethernet adapter. When you create an LHEA for a logical partition, you specify the resources that the logical partition can use on the actual physical HEA. Each logical partition can have one LHEA for each physical HEA on the managed system. Each LHEA can have one or more logical ports, and each logical port can connect to a physical port on the HEA.

You can create an LHEA for a logical partition using either of the following methods:

- You can add the LHEA to a partition profile, shut down the logical partition, and reactivate the logical partition using the partition profile with the LHEA.
- You can add the LHEA to a running logical partition using dynamic partitioning. This method can be used for Linux logical partitions only if you install the following operating systems on the logical partition:
 - Red Hat® Enterprise Linux version 4.6, or later
 - Red Hat Enterprise Linux version 5.1, or later
 - SUSE Linux Enterprise Server Version 10, or later
 - SUSE Linux Enterprise Server Version 11, or later

When you activate a logical partition, the LHEAs in the partition profile are considered to be required resources. If the physical HEA resources required by the LHEAs are not available, then the logical partition cannot be activated. However, when the logical partition is active, you can remove any LHEAs you want from the logical partition. For every active LHEA that you assign to an IBM i logical partition, IBM i requires 40 MB of memory.

After you create an LHEA for a logical partition, a network device is created in the logical partition. This network device is named `entX` on AIX logical partitions, `CMNXX` on IBM i logical partitions, and `ethX` on Linux logical partitions, where `X` represents sequentially assigned numbers. The user can then set up TCP/IP configuration like a physical Ethernet device to communicate with other logical partitions.

You can configure a logical partition so that it is the only logical partition that can access a physical port of an HEA by specifying *promiscuous mode* for an LHEA that is assigned to the logical partition. When an LHEA is in promiscuous mode, no other logical partitions can access the logical ports of the physical port that is associated with the LHEA that is in promiscuous mode. You might want to configure a logical partition to promiscuous mode in the following situations:

- If you want to connect more than 16 logical partitions to each other and to an external network through a physical port on an HEA, you can create a logical port on a Virtual I/O Server and configure an Ethernet bridge between the logical port and a virtual Ethernet adapter on a virtual LAN. This allows all logical partitions with virtual Ethernet adapters on the virtual LAN to communicate with the physical port through the Ethernet bridge. If you configure an Ethernet bridge between a logical port and a virtual Ethernet adapter, the physical port that is connected to the logical port must have the following properties:

- The physical port must be configured so that the Virtual I/O Server is the promiscuous mode logical partition for the physical port.
- The physical port can have only one logical port.
- You want the logical partition to have dedicated access to a physical port.
- You want to use tools such as tcpdump or iptrace.

A logical port can communicate with all other logical ports that are connected to the same physical port on the HEA. The physical port and its associated logical ports form a logical Ethernet network. Broadcast and multicast packets are distributed on this logical network as though it was a physical Ethernet network. You can connect up to 16 logical ports to a physical port using this logical network. By extension, you can connect up to 16 logical partitions to each other and to an external network through this logical network. The actual number of logical ports that you can connect to a physical port depends upon the Multi-Core Scaling value of the physical port group. It also depends on the number of logical ports that have been created for other physical ports within the physical port group. By default, the Multi-Core Scaling value of each physical port group is set to 4, which allows four logical ports to be connected to the physical ports in the physical port group. To allow up to 16 logical ports to be connected to the physical ports in the physical port group, you must change the Multi-Core Scaling value of the physical port group to 1 and restart the managed system.

You can set each logical port to restrict or allow packets that are tagged for specific VLANs. You can set a logical port to accept packets with any VLAN ID, or you can set a logical port to accept only the VLAN IDs that you specify. You can specify up to 20 individual VLAN IDs for each logical port.

The physical ports on an HEA are always configured on the managed system level. If you use an HMC to manage a system, you must use the HMC to configure the physical ports on any HEAs belonging to the managed system. Also, the physical port configuration applies to all logical partitions that use the physical port. (Some properties might require setup in the operating system as well. For example, the maximum packet size for a physical port on the HEA must be set on the managed system level using the HMC. However, you must also set the maximum packet size for each logical port within the operating system.) By contrast, if a system is unpartitioned and is not managed by an HMC, you can configure the physical ports on an HEA within the operating system as if the physical ports were ports on a regular physical Ethernet adapter.

HEA hardware does not support half-duplex mode.

You can change the properties of a logical port on an LHEA by using dynamic partitioning to remove the logical port from the logical partition. You can also add the logical port back to the logical partition using the changed properties. If the operating system of the logical partition does not support dynamic partitioning for LHEAs, and you want to change any logical port property other than the VLANs on which the logical port participates, you must set a partition profile for the logical partition so that the partition profile contains the wanted logical port properties, shut down the logical partition, and activate the logical partition using the new or changed partition profile. If the operating system of the logical partition does not support dynamic partitioning for LHEAs, and you want to change the VLANs on which the logical port participates, you must remove the logical port from a partition profile belonging to the logical partition, shut down and activate the logical partition using the changed partition profile, add the logical port back to the partition profile using the changed VLAN configuration, and shut down and activate the logical partition again using the changed partition profile.

Related concepts

Virtual Ethernet

Virtual Ethernet allows logical partitions to communicate with each other without having to assign physical hardware to the logical partitions.

Related information

Shared Ethernet Adapters

[Integrated Virtual Ethernet Adapter Technical Overview and Introduction](#)

Tagged resources for IBM i logical partitions

When you create an IBM i logical partition using the Hardware Management Console (HMC), you must tag I/O adapters (IOAs) to perform specific functions for the IBM i logical partition.

A *tagged resource* is an IOA that is selected because it controls a device that performs a specific function for a logical partition. The HMC and the IBM i operating system use this tagging to locate and use the correct I/O device for each I/O function. For example, when you create an IBM i partition profile, you must tag the I/O device that you want the IBM i logical partition to use as its load source. The tag allows the HMC to locate the load source when you activate the logical partition using the partition profile.

You can tag the IOA that controls the I/O device that you want to use. Tagging the IOA allows you to specify the exact I/O device that you want to use.

The following table lists and describes the device types that are tagged and indicates whether you must tag the device type for IBM i logical partitions.

<i>Table 11. Devices associated with tagged IOAs</i>		
Device	Description	Tagging required for IBM i logical partitions?
Alternate restart device	This device can be a tape drive or an optical device. The media in the alternate restart device is what the system uses to start from when you perform a D-mode initial program load (IPL). The alternate restart device loads the Licensed Internal Code that is contained on the removable media instead of the code on the load source disk unit.	Yes
Logical partition console	The first workstation that the system activates in the logical partition and the only device it activates on a manual IPL. The logical partition assumes that a console is always available for use.	Yes (if you are using a console device other than the HMC)
Load source disk unit	Each IBM i logical partition must have 1 disk unit designated as the load source. The system uses the load source to start the logical partition. The system always identifies this disk unit as unit number 1.	Yes

If you use the Virtual Partition Manager to create logical partition on your managed system, you do not have to tag I/O devices for these I/O functions. The IBM i logical partition automatically owns all physical I/O resources on the managed system, and the Virtual Partition Manager automatically tags the I/O device to use for each I/O function. The Virtual Partition Manager tags I/O devices for I/O functions based upon server model and location within the server. If you are partitioning a new server using the Virtual Partition Manager and have ordered the server with preinstalled IBM i, then you do not have to verify the placement of I/O devices within your new server.

Related concepts

I/O device assignment in partition profiles

I/O devices are assigned to partition profiles either on a slot-by-slot basis, or on logical port basis in the case of shared mode single root I/O virtualization (SR-IOV) adapters. For I/O devices that are assigned

to partition profiles on a slot-by-slot basis, most I/O devices can be assigned to a partition profile on the HMC as required or as allocated. For SR-IOV logical ports, I/O devices are always assigned to a profile as required.

Load source placement rules for IBM i logical partitions

You must properly place a disk unit within a system unit or expansion unit before you use the disk unit as the load source for an IBM i logical partition. The placement rules depend upon the server unit or expansion unit in which the load source is located and, sometimes on, the I/O adapter (IOA) that controls the load source.

Note: The information provided does not replace the System Planning Tool (SPT). Use this information as a resource with the SPT output. Its purpose is to assist you in the load source placement for your IBM i logical partitions.

- There is no specific slot requirement for serial-attached SCSI (SAS) disk units. Any slot can contain the load source.
- The load-source IOA must be specified when you create your logical partition.
- Disk compression must be disabled for the load source disk.
- Disk units must have at least 17 GB of usable capacity.
- Disk mirroring requires two load source disk devices in valid load-source positions.
- Any disk IOA that can attach to a system capable of having logical partitions can be used for additional storage capacity after the special requirements for the load-source disk are met.
- Each logical partition has its own single-level storage and hence its own ASP configuration. The ASP configuration rules that apply to systems without logical partitions also apply to logical partitions.
- Disk protection can be defined for a logical partition in the same way as for a nonpartitioned system: parity protection (RAID), mirroring, or mixed. Bus-level mirroring requires two buses in the logical partition.
- Disk units that are already in use by a logical partition cannot be easily added to a different logical partition. You must first remove them from the configuration of the logical partition that is using the disk units before you add them to a different logical partition. In doing so, the system automatically moves any user or system data to other disk units in the same ASP.

Alternate restart device placement rules for IBM i logical partitions

You can use the internal optical device in the system unit, or you can use an external tape or optical device, to load the Licensed Internal Code and IBM i to the load source disk unit of an IBM i logical partition.

The only supported internal device for alternate restart is the Slimline DVD drive in the removable media slot in the system unit.

IBM i logical partitions have the following rules for the external alternate restart devices:

- The alternate restart device must be connected to bus 0 or port 0 of the IOA.
- The alternate restart IOA is specified during logical partition setup.

Switchable devices for IBM i logical partitions

When you set up an I/O adapter (IOA) so that it can be switched from one logical partition to another, you can share the devices that are associated with that IOA among many IBM i logical partitions.

When you switch an IOA, you take the control of the devices away from one logical partition and give it to another without restarting the server or the logical partition. Before switching the IOA to another logical partition, you must ensure that the device is not in use.

IOAs that are good candidates for switching between logical partitions include IOAs that are attached to high-cost devices or low-use or low-demand devices.



Attention: When switching IOAs that control disk units, ensure that all disk units that belong to that specific IOA are first removed from the auxiliary storage pool and are in a unconfigured status.

Virtual OptiConnect for IBM i logical partitions

The virtual OptiConnect feature provides high-speed interpartition communication within a managed system. The Virtual OptiConnect feature emulates external OptiConnect hardware by providing a virtual bus between logical partitions.

The virtual OptiConnect feature can be used only for communications between IBM i logical partitions. If you must enable communications with AIX or Linux logical partitions, use virtual Ethernet instead of the virtual OptiConnect feature.

To use the virtual OptiConnect feature on a logical partition, you must install OptiConnect for IBM i (a priced optional feature) on each IBM i logical partition that is to use virtual OptiConnect. If you use the Hardware Management Console (HMC) to create logical partitions on your managed system, you must also check the partition profile properties for each IBM i logical partition that is to use virtual OptiConnect and ensure that the **Use virtual OptiConnect** option is selected on the **OptiConnect** tab.

You can use the virtual OptiConnect feature without any additional hardware requirements.

Related information

[OptiConnect](#)

Expansion unit

You can add expansion units to many of the models to support additional features and devices. If you want to create logical partitions on your server, you must add an expansion unit that contains the additional hardware that you need for each logical partition.

Some expansion units can support only disk units (storage expansion unit), while others can support various hardware (system expansion unit). Expansion units generally contain one or more system I/O buses with various I/O devices.

If you assign more than 144 I/O slots to an AIX or Linux logical partition, ensure that the boot device for the logical partition is within the first 144 slots that are assigned to the logical partition. Also ensure that any PCIe3 2-Port 40GbE NIC RoCE QSFP+ Adapter (FC EC3A or FC EC3B) is within the first 144 slots. You can view the devices that are assigned to the first 144 slots of a logical partition by viewing the partition properties of the logical partition. Select the **Hardware** tab, then select the **I/O** tab, and then click the **Bus** column of the table to sort the devices in ascending order.

5250 CPW for IBM i logical partitions

5250 commercial processing workload (5250 CPW) is the capacity to perform 5250 online transaction processing (5250 OLTP) tasks on IBM i logical partitions.

A *5250 OLTP task* is a task that uses the 5250 data stream. Examples of 5250 OLTP tasks include the following:

- Any form of 5250 emulation, including Hardware Management Console (HMC) 5250, IBM Host On-Demand, IBM Personal Communications, and the 5250 emulation in the IBM i Access for Windows, Web, and Linux products
- 5250 Telnet or 5250 Display Station Pass-Through (DSPT) workstations
- Screen scrapers
- Interactive system monitors

You can use the IBM WebFacing tool to convert your 5250 OLTP applications into web-based applications that no longer need to use the 5250 data stream.

Application support for Linux logical partitions

Learn how to integrate Linux with IBM i applications and data.

Samba support with IBM i NetServer

Server Message Block (SMB) is a file-sharing protocol that is commonly used by Windows PCs. Whenever a network drive is mapped from a Windows PC to another Windows PC, the SMB TCP/IP protocol is being used.

Samba implements the SMB/CIFS standard on UNIX operating systems. This protocol enables file sharing among SMB-enabled operating systems, including IBM i with NetServer.

Samba allows Linux PCs and servers to interact with existing Windows PCs and file servers without requiring any additional software. IBM i NetServer supports Linux Samba clients.

You can use a Samba server to run printers and authenticate users, share files, and directories, just like Microsoft Windows PCs. Samba can also act as a Primary Domain Controller (PDC) or as a Backup Domain Controller (BDC) in your Windows network. You can use it to run OpenLDAP and add LDAP function to your Windows Network without the expense. You can use Samba and NetServer to share printers and files on IBM Power Systems or Linux partitions.

Accessing IBM i data by using Linux ODBC driver

The IBM i Open Database Connectivity (ODBC) driver for Linux allows you to access the IBM i database data from Linux applications that are written to the ODBC API. It is based on the ODBC driver in the IBM i Access for Windows product.

Related information

[System i Access for Linux Open Database Connectivity](#)

Examples: Logically partitioned systems

You can use the logical partitioning examples to consolidate servers, use computing resources more efficiently, and increase the flexibility of your enterprise.

Creating multiple client environments

You provide high-availability e-commerce services to a number of clients. You provide computing resources, applications, and technical support to each client, and each client can independently configure and use the applications that are running on the computing resources that you provide. In such an environment, it is essential to isolate the clients so that the clients can access only their resources. However, dedicating a physical server to each client is cost prohibitive, and does not allow you to easily increase or decrease the amount of computing resources that are used by each client.

Therefore, you decide to create a logical partition for each client. You install an operating system and applications on each logical partition. You can then use dynamic partitioning to add resources to logical partitions or remove resources from logical partitions as needed. If a client stops using your service, you can delete the logical partition for that client and reassign the resources to other logical partitions.

Testing new applications

You are a furniture manufacturer that uses an application to track inventory at your plant. A new version of the application is now available. You want to test this new version before you use it on your production server, but you do not have any money to buy separate test hardware.

Therefore, you decide to create a separate test environment on your managed system. You remove resources from the existing production environment, and you create a new logical partition that contains the resources that you removed from the production environment. You install an operating system and the new version of the inventory application on the logical partition. You can then use dynamic partitioning to move resources from the test logical partition to the production logical partition during peak production

demand, and then return the resources to the test logical partition during testing. When you finish testing, you can delete the test logical partition, add the resources back to the production logical partition, and install the new version of the inventory application on the production system.

Integrating new acquisitions

You acquired a new company. Your new acquisition does not use the same applications for payroll, inventory, and billing that you do. You plan to consolidate your two companies onto a single set of applications, but it takes time to implement this consolidation. In the meantime, you are under pressure to reduce data center costs quickly.

Therefore, you decide to create logical partitions for the applications that are used by your new acquisition. You install an operating system and the applications that are used by the new company on the logical partition. If the combined workloads require more resources, you can use Capacity Upgrade on Demand (CUoD) to add processors and memory to the managed system, and then use dynamic partitioning to add these resources to the logical partitions. This solution allows you to save hardware costs immediately while you determine the best way to consolidate onto a single set of applications.

Planning for logical partitions

You can create logical partitions to distribute resources within a single server and make it function as if it were two or more independent servers. Before you create logical partitions, you must assess your current and future needs. You can then use this information to determine the hardware configuration that will meet your current needs and serve as a base for meeting your future needs.

Planning for logical partitions is a multi step process. The following tasks are recommended for planning for logical partitions.

_ **Assess your needs**

- Compile a list of the questions that you must answer before you create logical partition on an existing system or place your order for new hardware. The following is the list of questions:
 - What are your existing workloads? How many resources do these workloads currently require (during typical usage and at peak usage)?
 - What are your future needs? How will your existing workloads grow over the life of your system? How many new workloads do you have to support over the life of your system?
 - Do you have an existing system onto which you can consolidate the workloads? Must you upgrade the existing system before you consolidate the workloads? Does it make more sense to purchase a new system for these workloads?
 - What physical infrastructure will you have to support any new hardware? Can your current location accommodate the new hardware? Must you upgrade your power infrastructure or your cooling infrastructure?
 - Will your new hardware work with your existing hardware?
 - Which hardware features will you use? For example, do you want to use virtual I/O to consolidate I/O resources? Must you obtain activation codes or enablement codes to use these features?
 - Must you obtain more licenses to run your applications? If so, how many more licenses do you need?
 - Does the support strategy for your new hardware differ from the support strategy for your existing hardware? If so, what changes must you make to maximize the effectiveness of the new support strategy?
 - Must you migrate your workloads onto new hardware? If so, what must you do to migrate these workloads?

– **Learn about your system and its features**

- Your system has many features that allow you to use system resources more efficiently and simplify daily tasks. For more information about what these features are and how these features work, see [“Logical partition overview”](#) on page 2.

– **Learn about planning tools**

- IBM provides the following tools that you can use to assess your needs, determine the hardware that you need to accommodate existing and future needs, and compile an order for the hardware that you need:

IBM Prerequisite website

The [IBM Prerequisite website](#) provides you with compatibility information for hardware features. This site helps you plan a successful system upgrade by providing you with the prerequisite information for features that you currently have or plan to add to your system.

IBM Systems Workload Estimator

The [IBM Systems Workload Estimator \(WLE\)](#) estimates the computer resources that are required for Domino, WebSphere® Commerce, WebSphere, Web Serving, and traditional workloads. The WLE projects the most current server models that meet the capacity requirements that are within the CPU percent utilization objectives.

AIX Performance Toolbox for POWER®

The [AIX Performance Toolbox \(PTX\) for POWER](#) is a licensed program that provides a comprehensive tool for monitoring and tuning system performance in distributed environments.

– **Take inventory of your current environment**

- Monitor resource usage on your existing servers to determine the amounts of resources that you currently use in your operation. You will use this information as a basis for determining the resources that you require on the consolidated system. The Performance Monitor (PM) information that you gather from your existing systems gives you the information that you need to analyze existing workloads.

– **Perform capacity planning**

- Analyze the workloads that are to be consolidated onto your managed system and determine the amounts of resources that these workloads require. You will also want to calculate the resources that you will need for future growth and determine whether your hardware can accommodate this growth. To analyze your current workloads, use your PM information as input for the WLE. The WLE uses this input to determine the resources that you need for the consolidated workloads. The WLE also allows you to project how many resources you will need in the future.

– **Decide which tool you want to use to create logical partitions and manage the system**

- Determine whether you want to use the Hardware Management Console (HMC), or the Virtual Partition Manager to create logical partitions and manage the system. To learn about these tools, see [“Logical partitioning tools”](#) on page 10.

– **Decide if you want your operating systems to share I/O resources with each other**

- Determine whether you want to set your logical partitions to use virtual I/O resources from a Virtual I/O Server logical partition. For more information, see [Virtual I/O Server](#).

– **Design and validate your logical partition configuration**

- Design the logical partitions that you will create on the managed system, and assign resources to each logical partition so that the logical partitions can perform their assigned tasks efficiently.

- **Design network infrastructure to connect logical partitions with each other and with external networks**

Determine what types of physical and virtual adapters you want to use to connect logical partitions to each other and to external networks. For more information about the different methods that you can use to connect logical partitions with each other and with external networks, see [“I/O devices” on page 46](#).

- **Identify how the managed system communicates with the HMC**

- Determine how you want to connect your managed system and its logical partitions with the HMC that manages the system. For more information about the ways in which you can connect your managed system with the HMC, see [HMC network connections](#).

- **Determine a service and support strategy**

- Determine how to apply fixes to your server and identify problems that need to be reported to your service provider. The HMC can be configured to report most problems to your service provider automatically. For more information about how to set up the HMC to report problems, see [Configuring the local console to report errors to service and support](#).

- **Plan for software licensing in a partitioned environment**

- Determine the number of software licenses that you need for your logical partition configuration. For instructions, see [“Software licensing for IBM licensed programs on logical partitions” on page 75](#).

Trusted Boot

Trusted Boot is a feature of Power® Security and Compliance (PowerSC). Trusted Boot uses the Virtual Trusted Platform Module (VTPM) as described by the Trusted Computing Group. Up to 60 logical partitions per server can be configured to have their own unique VTPM by using the Hardware Management Console (HMC). The VTPM is used to record the system boot and in association with the AIX Trusted Execution technology, provides security and assurance of the boot image on disk, on the entire operating system, and in the application layers.

The VTPM is a software implementation of the Trusted Platform Module (TPM) specification, as described by the Trusted Computing Group. The Trusted Platform Module is implemented as a physical chip on computer systems.

You can create a VTPM as part of the initial logical partitioning (by using the HMC Partitioning wizard), or you can dynamically enable the device. When dynamically enabled, the VTPM becomes active only when the logical partition is restarted.

The VTPM enables the AIX environment of the logical partition to use Trusted Boot capability. When a VTPM is associated with a logical partition, being booted, components of the boot take cryptographic hashes of relevant data and of components that can be run in the future, for example the AIX boot loader. These cryptographic hashes are securely copied to storage that is controlled by the VTPM. After the logical partition is operational, other users can then securely retrieve the hashes by using a process that is known as remote attestation. The hashes can then be examined to determine whether the logical partition booted in a trusted configuration so that users can take action if required.

To use a VTPM, the logical partition must have the following resources:

- The maximum memory setting of the logical partition must be greater than of 1 GB for the active profile.
- Each VTPM requires permanent storage for the lifetime of the device. A normal logical partition uses 6 KB of system nonvolatile RAM. This storage requirement imposes a limitation on the number of VTPMs per server.

Permanent data that is stored by the VTPM contains sensitive information about the trust of the VTPM feature. For example, the first time each VTPM is operated a public-private key pair that is known as the Endorsement Key (EK) is generated and then permanently stored. This action allows the VTPM to

be identified by other users during the lifetime of the device. The permanent data, including the EK, is deleted when the VTPM device is removed by the console.

To maintain the sensitivity of the stored data, the data is secured by the trusted system key, which is under the control of the HMC. The trusted system key secures the VTPM data but has an impact on logical partition mobility, and suspend features for logical partitions that are enabled for VTPM. A logical partition that is enabled for VTPM must adhere to the following prerequisites to support the logical partition mobility, and suspend features:

- To migrate a logical partition with VTPM enabled, both systems must have the same trusted system key.
- To successfully change the trusted system key, no logical partition with VTPM enabled can be in the suspend state. The HMC cannot change the key until suspended logical partitions with VTPM enabled, are resumed or powered off.

Trusted Firewall

With Virtual I/O Server (VIOS) Version 2.2.1.4 or later, and POWER9 processor-based servers with firmware at level FW740, or later, you can use the Trusted Firewall feature. Trusted Firewall is a feature of the PowerSC Editions. You can use the Trusted Firewall feature to provide a virtual firewall that allows network filtering and control within the local server. The virtual firewall improves performance and reduces the consumption of network resources by allowing direct and secure network traffic between logical partitions that are on different VLANs of the same server.

With the Trusted Firewall feature, you can perform LAN routing functions between logical partitions on the same server by using the Security Virtual Machine (SVM) kernel extension. By using the Trusted Firewall feature, logical partitions that are on different virtual LANs of the same server can communicate by using the shared Ethernet adapter (SEA). Trusted Firewall is supported on AIX, IBM i, and Linux logical partitions.

Preparing to configure Active Memory Expansion

Before you configure Active Memory Expansion for a logical partition, you need to ensure that your system meets the configuration requirements. Optionally, you can run the Active Memory Expansion planning tool.

About this task

To prepare to configure Active Memory Expansion for a logical partition, complete the following tasks:

Procedure

1. Ensure that your system meets the following configuration requirements:
 - The server on which the logical partition runs is a POWER7, or later.
 - The AIX operating system that runs in the logical partition is at version 6.1 with Technology Level 4 and Service Pack 2, or later.
 - The HMC that manages the server is at version 7, release 7.1.0, or later.
2. Optional: Run the Active Memory Expansion planning tool, which is the **amepat** command, from the AIX command-line interface.

The planning tool monitors your current workload for a specified amount of time and generates a report. The report provides the following information:

- Several configuration possibilities for Active Memory Expansion on the logical partition.
- Recommendation for an initial configuration for Active Memory Expansion on the logical partition.

For each configuration possibility, and the recommended configuration, the planning tool provides the following configuration information:

- The amount of memory to assign to the logical partition.
- The amount of additional processing resources to assign to the logical partition.

- The expansion factor to set for the logical partition.
- The amount of memory that you save by configuring Active Memory Expansion on the logical partition. This statistic can help you determine whether Active Memory Expansion is right for your workload. Some workloads benefit more from Active Memory Expansion than others.

What to do next

After you prepare to configure Active Memory Expansion, you can enable Active Memory Expansion on the server by entering the activation code.

Configuration requirements for shared memory

Review the requirements for the system, Virtual I/O Server (VIOS), logical partitions, and paging space devices so that you can successfully configure shared memory.

System requirements

- The server must be a POWER7 processor-based server, or later.
- The server firmware must be at release 3.4.2, or later.
- The Hardware Management Console (HMC) must be at version 7 release 3.4.2, or later.
- The PowerVM Active Memory Sharing technology must be activated. The PowerVM Active Memory Sharing technology is available with the PowerVM Enterprise Edition for which you must obtain and enter a PowerVM Editions activation code. Only 512 byte block devices are supported for PowerVM Active Memory Sharing.

Paging VIOS partition requirements

- VIOS partitions that provide access to the paging space devices for the shared memory partitions that are assigned to the shared memory pool (hereafter referred to as *paging VIOS partitions*) cannot use shared memory. Paging VIOS partitions must use dedicated memory.
- Paging VIOS partitions must be at version 2.1.1, or later.
- On HMC-managed systems, consider configuring separate VIOS partitions as server partitions and paging VIOS partitions. For example, configure one VIOS partition to provide virtual resources to the shared memory partitions. Then, configure another VIOS partition as a paging VIOS partition.
- On HMC-managed systems, you can configure multiple VIOS partitions to provide access to paging space devices. However, you can assign only up to two of those VIOS partitions to the shared memory pool at any time.

Requirements for shared memory partitions

- Shared memory partitions must use shared processors.
- You can assign only virtual adapters to shared memory partitions. This means that you can dynamically add only virtual adapters to shared memory partitions. More specifically, the following table lists the virtual adapters that you can assign shared memory partitions.

Table 12. Virtual adapters that you can assign to shared memory partitions	
AIX and Linux shared memory partitions	IBM i shared memory partitions
– Virtual SCSI client adapters – Virtual Ethernet adapters – Virtual Fibre Channel client adapters – Virtual serial adapters	– Virtual SCSI client adapters – Virtual Ethernet adapters – Virtual Fibre Channel client adapters – Virtual serial server adapters

Table 13. Virtual adapters that you can assign to shared memory partitions

Linux shared memory partitions

- Virtual SCSI client adapters
- Virtual Ethernet adapters
- Virtual Fibre Channel client adapters
- Virtual serial adapters

You cannot assign Host Ethernet Adapters (HEA) or host connection adapters (HCA) to shared memory partitions.

- Shared memory partitions cannot use the barrier synchronization register.
- Shared memory partitions cannot use huge pages.
- AIX must be at version 6.1 Technology Level 3, or later, to run in a shared memory partition.
- IBM i must be at 6.1 with PTF SI32798, or later, to run in a shared memory partition.
- Virtual OptiConnect must not be enabled on IBM i shared memory partitions.
- SUSE Linux Enterprise Server must be at version 11, or later, to run in a shared memory partition.
- Red Hat Enterprise Server Version 6, or later, to run in a shared memory partition.
- You cannot configure IBM i logical partitions that provide virtual resources to other logical partitions as shared memory partitions. Logical partitions that provide virtual resources to other logical partitions in a shared memory environment must be VIOS partitions.

Requirements for paging space devices

- The paging space devices for AIX or Linux shared memory partitions must be at least the size of the maximum logical memory of the shared memory partition.
- The paging space devices for IBM i shared memory partitions must be at least the size of the maximum logical memory of the shared memory partition plus 8 KB for every megabyte. For example, if the maximum logical memory of the shared memory partition is 16 GB, its paging space device must be at least 16.125 GB.
- Paging space devices can be assigned only to one shared memory pool at a time. You cannot assign the same paging space device to a shared memory pool on one system and to another shared memory pool on another system at the same time.
- Paging space devices that are accessed by a single paging VIOS partition must meet the following requirements:
 - They can be physical or logical volumes.
 - They can be located in physical storage on the server or on a storage area network (SAN).
- Paging space devices that are accessed redundantly by two paging VIOS partitions must meet the following requirements:
 - They must be physical volumes.
 - They must be located on a SAN.
 - They must be configured with global IDs.
 - They must be accessible to both paging VIOS partitions.
 - The reserve attribute must be set to no reserve. (The VIOS automatically sets the reserve attribute to no reserve when you add the paging space device to the shared memory pool.)
- Physical volumes that are configured as paging space devices cannot belong to a volume group, such as the rootvg volume group.
- Logical volumes that are configured as paging space devices must be located in a volume group that is dedicated for paging space devices.

- Paging space devices must be available. You cannot use the physical volume or logical volume as a paging space device if it is already configured as a paging space device or virtual disk for another logical partition.
- Paging space devices cannot be used to boot a logical partition.
- After you assign a paging space device to the shared memory pool, you must manage the device by using the **Create/Modify Shared Memory Pool** wizard on the HMC. Do not change or remove the device by using other management tools.

Related concepts

[Paging space devices on systems that are managed by an HMC](#)

Learn about the location requirements, size requirements, and redundancy preferences for paging space devices on systems that are managed by a Hardware Management Console (HMC).

Configuration requirements and restrictions for the remote restart capability of a logical partition

An AIX, Linux, or IBM i logical partition that supports the remote restart feature must have its configuration information and persistent data stored external to the server on persistent storage.

With the Hardware Management Console (HMC) Version 7.6.0, or later, a logical partition can be enabled for remote restart on any server that supports the remote restart capability. To perform a remote restart operation by using the HMC command line, the HMC must be at version 8.1.0 or later. A logical partition can recover from a server outage by starting on another server.

The following are the configuration requirements for a logical partition with the remote restart capability:

- When the HMC is at Version 8.2.0, or later, and when you select the PowerVM simplified remote restart feature, you need not assign a reserved storage device from the reserved storage device pool to the logical partition before the logical partition is activated.
- When the Hardware Management Console (HMC) is at Version 8.4.0, or later, and the Virtual I/O Server (VIOS) at Version 2.2.4.0, or later, the remote restart of logical partitions that use shared storage pool devices is supported.
- When the HMC is at Version 9.1.940, or later, and when the firmware is at level FW940, or later, you can remotely restart a logical partition that has migratable single root I/O virtualization (SR-IOV) logical ports.

Note: When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

- When the HMC is at Version 9.2.950, or later, and the firmware is at a level FW950, or later, for the remote restart operation to be successful on the logical partition that supports platform keystore capability, the user-defined system key that is configured on both the source and the destination system must match.

Restriction:

The following are the restrictions for a logical partition that supports the remote restart feature:

- The logical partition must not have physical I/O adapters.
- If you are using HMC Version 9.1.930, or earlier, the logical partition must not have SR-IOV logical ports.
- The logical partition must not be a full system partition, or a Virtual I/O Server (VIOS) partition.
- The logical partition must not be enabled for redundant error path reporting.
- The logical partition must not have a barrier-synchronization register (BSR).
- The logical partition must not have huge pages.

- The logical partition must not have its rootvg volume group on a logical volume or have any exported optical devices.
- The logical partition must not be set to automatically start with the server.
- The logical partition must not have any Host Channel Adapter resources.
- The logical partition must not be set as the service partition for the server.
- The logical partition must not have any Host Ethernet Adapter resources.
- The logical partition must not belong to a workload group.
- The logical partition must not use shared memory.
- The logical partition must not have a Virtual Trusted Platform Module (VTPM) enabled.
- The logical partition must not use Virtual Station Interface.

In addition to the restrictions for a logical partition that supports the remote restart feature, IBM i partitions have some additional restrictions.

Restriction:

The following list illustrates the restrictions for an IBM i logical partition that supports the remote restart feature:

- The logical partition must not have a virtual server SCSI adapter. Also, the logical partition must not have a virtual SCSI client adapter that is associated with a server adapter that is not on a VIOS partition.
- The logical partition must not have HSL (High-Speed Link) OptiConnect or Virtual OptiConnect enabled.

Verifying that the server supports partitions that are capable of the simplified version of the remote restart feature

Before planning to enable the simplified version of the remote restart capability of a logical partition, verify that the server supports partitions that are capable of the simplified version of the remote restart feature by using the Hardware Management Console (HMC).

About this task

To verify that the server supports partitions that are capable of the simplified version of the remote restart feature, complete the following steps:



1. In the navigation pane, click the **Resources** icon.
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **Actions > View System Properties**. The **Properties** page is displayed.
4. Click **Licensed Capabilities**. The Licensed Capabilities page lists the features that are supported by the server.
5. In the Licensed Capabilities page, verify the list of features displayed.
 - If **PowerVM Partition Simplified Remote Restart Capable** is marked by the  icon, server supports the simplified version of the remote restart feature.
 - If **PowerVM Partition Simplified Remote Restart Capable** is marked by the  icon, the server does not support the simplified version of the remote restart feature.
6. Click **OK**.

Verifying that the logical partition supports the simplified version of the remote restart feature

You can use the Hardware Management Console (HMC) to verify whether the logical partition supports the simplified version of the remote restart feature.

About this task

To verify that the logical partition supports the simplified version of the remote restart feature, complete the following steps:



1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions > View Partition Properties**.
4. Click the **General Properties** tab.
 - If the **Simplified Remote Restart** check box is selected, the logical partition supports the simplified version of the remote restart feature.
 - If the **Simplified Remote Restart** check box is not selected, the logical partition does not support the simplified version of the remote restart feature.
5. Click **OK**.

Verifying that the server supports Virtual Trusted Platform Module

Before planning to enable a Virtual Trusted Platform Module (VTPM) on a logical partition, verify that the server supports VTPM by using the Hardware Management Console (HMC).

About this task

To verify that the server supports partitions that are capable of VTPM, complete the following steps:



1. In the navigation pane, click the **Resources** icon .
2. In the work pane, select the system and click **Actions > View System Properties**. The **Properties** page is displayed.
3. Click **Advanced**. The server supports VTPM if you can view information about VTPM.
4. Click **OK**.

Enabling the platform keystore capability on a logical partition

With the HMC version 9.2.950, you can enable the platform key store capability if the system supports the feature.

Before enabling the platform keystore feature on a logical partition, verify that the IBM Power system server supports the platform keystore capability by using the **lssyscfg** command along with **-F** attribute. Before running this command, ensure that the system is in a stand by or operating state. If the system supports the platform keystore capability, you can also use the **lssyscfg** command to view the range of the keystore size.

When the HMC is at V9.2.950.0, or later, and when the firmware is at level FW950, or later, you can choose the keystore size for a logical partition. You can choose the value for the keystore size within the range supported by the system. The value of 0 kilobytes (KB) indicates that the keystore function is disabled for the logical partition. To view the minimum and maximum value for the keystore size, type the following the command:

```
lssyscfg -r sys -m <system_name> -Flpar_keystore_min_kbytes, lpar_keystore_max_kbytes
```

Related information

[lssyscfg command](#)

[chtskey command](#)

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 **lldp** 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.

Verifying that the server uses the virtual server network

Before planning to enable the virtual server network (VSN), verify that the server uses VSN by using the Hardware Management Console (HMC).

About this 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 logical partitions. The VEPA switching mode uses the features that are supported by the advanced virtual Ethernet switch technology. A logical partition whose virtual Ethernet adapters use virtual switches that are enabled with the VEPA switching mode, uses VSN.

You can use the **lssyscfg** command to verify that the server uses VSN.

Verifying that the server supports single root I/O virtualization

Before you enable single root I/O virtualization (SR-IOV) shared mode for an SR-IOV capable adapter, verify that the server supports the SR-IOV feature by using the Hardware Management Console (HMC). 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.

About this task

To verify that the server supports SR-IOV, complete the following steps:



1. In the navigation pane, click the **Resources** icon.
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **Actions** > **View System Properties**. The **Properties** page is displayed.
4. Click **Licensed Capabilities**. The Licensed Capabilities page lists the features that are supported by the server.
5. In the Licensed Capabilities page, verify the list of features displayed.
 - If **SR-IOV Capable** is marked by the  icon, the SR-IOV adapter can be configured in the shared mode and can be shared by multiple logical partitions.
 - If **SR-IOV Capable** is marked by the  icon, the SR-IOV adapter can be configured in the shared mode but can be used by only one logical partition.
 - If **SR-IOV Capable** is not displayed, the server does not support the SR-IOV feature.
6. Click **OK**.

Verifying the logical port limit and the owner of the SR-IOV adapter

You can view the logical port limit and the owner of the single root I/O virtualization (SR-IOV) adapter by using the Hardware Management Console (HMC).

About this task

To view the logical port limit and the owner of the SR-IOV adapter, complete the following steps:



1. In the navigation pane, click the **Resources** icon.
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **Actions** > **View System Properties**. The **Properties** page is displayed.
4. Click **Licensed Capabilities**. The Licensed Capabilities page lists the features that are supported by the server.
5. In the **Properties** area, click the **Processor, Memory, I/O** tab. In the **Physical I/O Adapters** area, the table displays the **SR-IOV capable (Logical Port Limit)** and the **Owner** details about the SR-IOV adapter.
 - The **SR-IOV capable (Logical Port Limit)** column displays whether the slot or the adapter is SR-IOV capable, and the maximum number of logical ports this slot or the adapter can support. If the slot or the adapter is SR-IOV capable but is assigned to a partition, the **SR-IOV capable (Logical Port Limit)** column indicates that the slot or the adapter is in the dedicated mode.

- The **Owner** column displays the name of the current owner the physical I/O. The value of this column can be any of the following values:
 - When an SR-IOV adapter is in the shared mode, **Hypervisor** is displayed in this column.
 - When an SR-IOV adapter is in the dedicated mode, **Unassigned** is displayed when the adapter is not assigned to any partition as a dedicated physical I/O.
 - When an SR-IOV adapter is in the dedicated mode, the logical partition name is displayed when the adapter is assigned to any logical partition as a dedicated physical I/O.

Verifying that the server supports IBM i native I/O capability

You can verify that the server supports IBM i native I/O capability by running the **lssyscfg** command that is available in the Hardware Management Console (HMC) command-line interface.

Procedure

To verify whether the server supports IBM i native I/O capability, type the following command from the HMC command line:

```
lssyscfg -r sys -m <server> -F os400_native_io_capable
```

where *server* is the user-defined name of the server you want to check. The *os400_native_io_capable* attribute can have any one of the following values:

- unavailable - When the server firmware is at a level that is earlier than FW860. When this value is returned, the HMC does not know whether the server supports IBM i native I/O capability.
- 0 - When the server does not support IBM i native I/O capability.
- 1 - When the server supports IBM i native I/O capability.

Preparing to configure shared memory

Before you configure the shared memory pool and create logical partitions that use shared memory (hereafter referred to as *shared memory partitions*), you need to plan for the shared memory pool, the shared memory partitions, the paging space devices, and the Virtual I/O Server logical partitions (hereafter referred to as *paging VIOS partitions*).

Related concepts

Paging space devices on systems that are managed by an HMC

Learn about the location requirements, size requirements, and redundancy preferences for paging space devices on systems that are managed by a Hardware Management Console (HMC).

Logical memory

Logical memory is the address space, assigned to a logical partition, that the operating system perceives as its main storage. For a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*), a subset of the logical memory is backed up by physical main storage and the remaining logical memory is kept in auxiliary storage.

Preparing to configure shared memory on a system that is managed by an HMC

Before you configure the shared memory pool and create logical partitions that use shared memory, you need to determine the size of the shared memory pool, the amount of memory to assign to each shared memory partition, the number of paging space devices to assign to the shared memory pool, and the redundancy configuration of the Virtual I/O Server logical partitions that you assign to the shared memory pool.

Before you begin

Before you start, verify that your system meets the requirements for configuring shared memory. For instructions, see [“Configuration requirements for shared memory”](#) on page 65.

About this task

To prepare to configure the shared memory pool and shared memory partitions, complete the following steps:

Procedure

1. Assess your needs, take inventory of your current environment, and plan for capacity. For instructions, see [“Planning for logical partitions” on page 61](#). Determine the following information:
 - a) Determine the number of shared memory partitions to assign to the shared memory pool.
 - b) Determine the amount of logical memory to assign as the allocated, minimum, and maximum logical memory for each shared memory partition.

You can apply the same general guidelines that you might use to assign the allocated, minimum, and maximum dedicated memory to logical partitions that use dedicated memory. For example:

 - Do not assign the maximum logical memory to a value that is higher than the amount of logical memory that you plan to dynamically add to the shared memory partition.
 - Set the minimum logical memory to a value that is high enough for the shared memory partition to successfully activate.
2. Determine the amount of physical memory to assign to the shared memory pool.

For instructions, see [“Determining the size of the shared memory pool” on page 74](#).
3. Prepare for paging space devices:
 - a) Determine the number of paging space devices to assign to the shared memory pool.

The HMC assigns one paging space device to each shared memory partition that is active. Thus, the fewest number of paging space devices that must be assigned to the shared memory pool is equal to the number of shared memory partitions that you plan to run simultaneously. For example, you assign 10 shared memory partitions to the shared memory pool and you plan to run eight of the shared memory partitions simultaneously. Thus, you assign at least eight paging space devices to the shared memory pool.
 - b) Determine the size of each paging space device:
 - For AIX and Linux shared memory partitions, the paging space device must be at least the size of the maximum logical memory size of the shared memory partition that you identified in step 1b. For example, you plan to create an AIX shared memory partition with a maximum logical memory size of 16 GB. The paging space device must be at least 16 GB.
 - For IBM i shared memory partitions, the paging space device must be the size of the maximum logical memory size of the shared memory partition that you identified in step 1b multiplied by 129/128. For example, you plan to create an IBM i shared memory partition with a maximum logical memory size of 16 GB. The paging space device must be at least 16.125 GB.
 - Consider creating paging space devices that are large enough to be used by shared memory partitions with multiple partition profiles.
 - c) Determine whether each paging space device resides in physical storage on the server or on a storage area network (SAN).

Paging space devices that are accessed by a single Virtual I/O Server (VIOS) logical partition can be located in physical storage on the server or on a SAN. Paging space devices that are redundantly accessed by two paging VIOS partitions must be located on a SAN.
4. Prepare for paging VIOS partitions:
 - a) Determine which Virtual I/O Server (VIOS) logical partitions can be assigned to the shared memory pool as paging VIOS partitions.

A paging VIOS partition provides access to the paging space devices for the shared memory partitions that are assigned to the shared memory pool. A paging VIOS partition can be any active Virtual I/O Server (version 2.1.1, or later) that has access to the paging space devices that you plan to assign to the shared memory pool.
 - b) Determine the number of paging VIOS partitions to assign to the shared memory pool.

You can assign 1 or 2 paging VIOS partitions to the shared memory pool:

- When you assign a single paging VIOS partition to the shared memory pool, it must have access to all of the paging space devices that you plan to assign to the shared memory pool.
 - When you assign two paging VIOS partitions to the shared memory pool, each paging space device that you plan to assign to the shared memory pool must be accessible to at least one paging VIOS partition. However, usually when you assign two paging VIOS partitions to the shared memory pool, they redundantly access one or more paging space devices.
- c) If you plan to assign two paging VIOS partitions to the shared memory pool, determine how you want to configure redundancy for the shared memory partitions:
- i) Determine which shared memory partitions to configure to use redundant paging VIOS partitions. For each shared memory partition, this means that both paging VIOS partitions can access the shared memory partition's paging space device.
 - ii) Determine which paging VIOS partition to assign as the primary paging VIOS partition and which paging VIOS partition to assign as the secondary paging VIOS partition for each shared memory partition. The hypervisor uses the primary paging VIOS partition to access the paging space device that is assigned to the shared memory partition. If the primary VIOS partition becomes unavailable, the hypervisor uses the secondary paging VIOS partition to access the paging space device that is assigned to the shared memory partition.
5. Determine the number of additional processor resources that are needed for the paging VIOS partitions.

To read and write data between the paging space devices and the shared memory pool, the paging VIOS partitions require more processing resources. The amount of additional processing resources that are needed depends on the frequency that the paging VIOS partition reads and writes the data. The more frequently that the paging VIOS partition reads and writes the data, the more frequently the paging VIOS partition performs I/O operations. More I/O operations require more processing power. In general, the frequency that the paging VIOS partition reads and writes data can be affected by the following factors:

- The degree to which the shared memory partitions are overcommitted. In general, greatly overcommitted shared memory partitions require the paging VIOS partition to read and write data more often than slightly overcommitted shared memory partitions.
- The I/O rates of the storage subsystem on which the paging space devices are located. In general, paging space devices with faster I/O rates (such as a SAN) enable the paging VIOS partition to read and write data more often than paging space devices with slower I/O rates (such as storage in the server).

You can use the [IBM Systems Workload Estimator \(WLE\)](#) to determine the number of processor resources that are needed for paging VIOS partitions.

Related concepts

Performance considerations for shared memory partitions

You can learn about performance factors (such as shared memory overcommitment) that influence the performance of a logical partition that uses shared memory (also referred to as a *shared memory partition*). You can also use shared memory statistics to help you determine how to adjust the configuration of a shared memory partition to improve its performance.

Determining the size of the shared memory pool

You need to consider the degree to which you want to over-commit the physical memory in the shared memory pool, the performance of the workloads when they are running in a shared memory configuration that is over-committed, and the minimum and maximum boundaries of the shared memory pool.

About this task

To determine the size of the shared memory pool, consider the following factors:

Procedure

1. Consider the degree to which you want to over-commit the physical memory in the shared memory pool.
 - When the sum of the physical memory that is currently used by the shared memory partitions is less than or equal to the amount of memory in the shared memory pool, the memory configuration is *logically over committed*. In a logically over committed memory configuration, the shared memory pool has enough physical memory to contain the memory that is used by all shared memory partitions at one point in time.
 - When the sum of the physical memory that is currently used by the shared memory partitions is greater than the amount of memory in the shared memory pool, the memory configuration is *physically over committed*. In a physically over committed memory configuration, the shared memory pool does not have enough physical memory to contain the memory that is used by all the shared memory partitions at one point in time. The hypervisor stores the difference of the physical and shared memory in the auxiliary storage.
2. Consider the performance of the workloads when running in a shared memory configuration that is over-committed. Some workloads perform well in a shared memory configuration that is logically over-committed, and some workloads can perform well in a shared memory configuration that is physically over-committed.

Tip: In general, more workloads perform better in logically over-committed configurations than physically over-committed configurations. Consider limiting the degree to which you physically over-commit the shared memory pool.

3. The shared memory pool must be large enough to meet the following requirements:
 - a) The shared memory pool must be large enough to provide each shared memory partition with its I/O entitled memory when all of the shared memory partitions are active. When you create a shared memory partition, the Hardware Management Console (HMC) automatically determines the I/O entitled memory for the shared memory partition.

After you activate the shared memory partitions, you can view statistics about how the operating systems use their I/O entitled memory and adjust the I/O entitled memory of the shared memory partitions accordingly.
 - b) A small portion of the physical memory in the shared memory pool is reserved for the hypervisor so that it can manage shared memory resources. The hypervisor requirement is a small amount of physical memory per shared memory partition + 256 MB.

Tip: To ensure that you can successfully activate the shared memory partitions, assign at least the following amount of physical memory to the shared memory pool: (the sum of the minimum logical memory that is assigned to all of the shared memory partitions that you plan to run concurrently) + (the required 256 MB of reserved firmware memory).

4. When the shared memory pool is equal to or greater than the sum of the assigned logical memory of all the shared memory partitions plus the required amount of reserved firmware memory, the initial shared memory configuration is not over-committed. Therefore, the amount of physical memory that you assign to the shared memory pool need not exceed the sum of the assigned logical memory of all the shared memory partitions plus the required amount of reserved firmware memory.

Software licensing for IBM licensed programs on logical partitions

If you use IBM licensed programs such as AIX and IBM i on a server with logical partitions, consider how many software licenses are required for your logical partition configuration. Careful consideration of your software might help minimize the number of software licenses that you must purchase.

Software license behavior varies by software product. Each solution provider has its own licensing strategy. If you use licensed programs from solution providers other than IBM, consult the documentation from those solution providers to determine the licensing requirements for those licensed programs.

With some servers, you can purchase IBM i licenses on a per-user basis. For more information about IBM i licenses, see [Working with software agreements and licenses](#) in the IBM Knowledge Center.

Many IBM licensed programs allow you to purchase licenses based on the number of processors that the licensed program uses on a managed system as a whole. An advantage of this processor-based licensing method is that it allows you to create multiple logical partitions without having to purchase separate licenses for each logical partition. Also, this method caps the number of licenses that you need for a managed system. You need never obtain more licenses for a single licensed program than the number of processors on the managed system.

The main complicating factor in calculating the number of licenses that are required on a managed system with logical partitions that use processor-based licensing is the fact that a logical partition that uses uncapped shared processors can use up to its assigned number of virtual processors. When you use processor-based licensing, ensure that the number of virtual processors on uncapped logical partitions are set so that each IBM licensed program does not use more processors than the number of processor-based licenses that you purchased for that IBM licensed program.

The number of licenses that are required for a single IBM licensed program on a managed system that uses processor-based licensing is the **lesser** of the following two values:

- The total number of activated processors on the managed system.
- The maximum number of processors that can be used by the IBM licensed program on the managed system. The maximum number of processors that can be used by the IBM licensed program on the managed system is the **sum** of the following two values:
 - The total number of processors that are assigned to all logical partitions that use dedicated processors and that run the IBM licensed program.
 - The sum of the maximum number of processing units that can run the IBM licensed program on **each** shared processor pool, rounded up to the next integer. The maximum number of processing units that can run the IBM licensed program on each shared processor pool is the **lesser** of the following two values:
 - The total number of processing units that are assigned to capped logical partitions that run the IBM licensed program, plus the total number of virtual processors that are assigned to uncapped logical partitions that run the IBM licensed program.
 - The maximum number of processing units that are specified for the shared processor pool. (For the default shared processor pool, this number is the total number of activated processors on the managed system minus the total number of processors that are assigned to all logical partitions that use dedicated processors and that are not set to share processors with shared processor logical partitions. Use of Capacity on Demand (CoD) can increase the number of activated processors on the managed system, which can cause the managed system to go out of compliance if you do not allow for CoD use. Also, if there are logical partitions that use dedicated processors, that run the IBM licensed program, and that are set to share processors with shared processor logical partitions, then you can deduct the processors for these dedicated processor logical partitions from the maximum number of processing units for the default shared processor pool total because you already counted these dedicated processors in the dedicated processor logical partition total.)

When you use processor-based licensing, ensure that the managed system is in compliance with the license agreement for each IBM licensed program that is installed on the managed system. If you have a managed system that can use multiple shared processor pools, you can use the multiple shared processor pool feature of the Hardware Management Console (HMC) to ensure that your managed system remains in compliance with these license agreements. You can configure a shared processor pool with a maximum processing unit value equal to the number of licenses that you have for your managed system, and then set all logical partitions that use the IBM licensed program so that they use that shared processor pool. The logical partitions in the shared processor pool cannot use more processors than the maximum processing unit value that is set for the shared processor pool, so the managed system remains in compliance with the per-processor license agreement.

For example, Company Y has obtained three processor-based IBM i licenses for a managed system with four processors and four logical partitions. The managed system has only one shared processing pool, and all four logical partitions use the shared processor pool, so all four of the processors on the managed system are in the shared processor pool. The configuration of the logical partitions is as follows.

Table 14. Logical partition configuration in compliance with license agreement

Logical partition name	Operating system	Processing mode	Sharing mode	Processing units	Virtual processors	Maximum number of processors that can be used by the logical partition
Partition A	IBM i	Shared	Uncapped	1.75	2	2.00 (the number of virtual processors for the uncapped shared logical partition)
Partition B	IBM i	Shared	Capped	0.60	1	0.60 (the number of processing units for the capped shared logical partition)
Partition C	IBM i	Shared	Capped	0.40	1	0.40 (the number of processing units for the capped shared logical partition)
Partition D	Linux	Shared	Uncapped	1.25	2	2.00 (the number of virtual processors for the uncapped shared logical partition)

This configuration has three IBM i logical partitions and one Linux logical partition on the managed system. The three IBM i logical partitions can use a maximum of 3.00 processors (2.00 for Partition A, 0.60 for Partition B, and 0.40 for Partition C). The managed system has three IBM i licenses, so the managed system is in compliance with the IBM i license agreement.

For an example of a logical partition configuration that is out of compliance with a licensing agreement, the system administrator at Company Y changes the sharing mode of Partition B and Partition C from capped to uncapped. The following table shows the new logical partition configuration.

Table 15. Logical partition configuration out of compliance with license agreement (first example)

Logical partition name	Operating system	Processing mode	Sharing mode	Processing units	Virtual processors	Maximum number of processors that can be used by the logical partition
Partition A	IBM i	Shared	Uncapped	1.75	2	2.00 (the number of virtual processors for the uncapped shared logical partition)
Partition B	IBM i	Shared	Uncapped	0.60	1	1.00 (the number of virtual processors for the uncapped shared logical partition)
Partition C	IBM i	Shared	Uncapped	0.40	1	1.00 (the number of virtual processors for the uncapped shared logical partition)
Partition D	Linux	Shared	Uncapped	1.25	2	2.00 (the number of virtual processors for the uncapped shared logical partition)

In this configuration, the three IBM i logical partitions can use a maximum of 4.00 processors (2.00 for Partition A, 1.00 for Partition B, and 1.00 for Partition C). The managed system has only three IBM i licenses, but requires a total of four IBM i licenses, so the managed system is out of compliance with the IBM i license agreement.

If you have a managed system that can use multiple shared processor pools, you can use the Hardware Management Console (HMC) to configure a shared processor pool with a maximum processing unit value of 3.00, and assign Partition A, Partition B, and Partition C to that shared processor pool. If you do this, Partition A, Partition B, and Partition C can continue to be uncapped. You would remain in compliance with the IBM i license agreement because the maximum processing unit value would ensure that IBM i uses no more than three processing units.

For another example of a logical partition configuration that is out of compliance with a licensing agreement, the system administrator at Company Y changes the sharing mode of Partition B and Partition C back to capped. However, the system administrator then moves 0.50 processing units from Partition D to Partition A. Before the system administrator is allowed to do this, the system administrator must increase the number of virtual processors on Partition A from 2 to 3. The following table shows the new logical partition configuration.

Table 16. Logical partition configuration out of compliance with license agreement (second example)

Logical partition name	Operating system	Processing mode	Sharing mode	Processing units	Virtual processors	Maximum number of processors that can be used by the logical partition
Partition A	IBM i	Shared	Uncapped	2.25	3	3.00 (the number of virtual processors for the uncapped shared logical partition)
Partition B	IBM i	Shared	Capped	0.60	1	0.60 (the number of processing units for the capped shared logical partition)
Partition C	IBM i	Shared	Capped	0.40	1	0.40 (the number of processing units for the capped shared logical partition)
Partition D	Linux	Shared	Uncapped	0.75	2	2.00 (the number of virtual processors for the uncapped shared logical partition)

In this configuration, the three IBM i logical partitions can use a maximum of 4.00 processors (3.00 for Partition A, 0.60 for Partition B, and 0.40 for Partition C). The managed system has only three IBM i licenses, but requires a total of four IBM i licenses, so the managed system is out of compliance with the IBM i license agreement.

Considerations other than licensed program agreements might constrain your ability to run IBM licensed programs on certain server models.

Related concepts

Processors

A *processor* is a device that processes programmed instructions. The more processors that you assign to a logical partition, the greater the number of concurrent operations that the logical partition can run at any given time.

Minimum hardware configuration requirements for logical partitions

Each logical partition requires at least a certain amount of hardware resources. You can assign hardware resources directly to a logical partition, or you can set the logical partition to use the hardware resources that are assigned to another logical partition. The minimum hardware configuration requirements for each logical partition depend on the operating system or software that is installed on the logical partition.

The following table lists the minimum hardware requirements for logical partitions.

Table 17. Minimum hardware requirements for logical partitions

Minimum requirement	AIX and Linux	IBM i
Processor	One dedicated processor or 0.1 processing unit, or 0.05 processing unit when the firmware is at level FW760, or later.	One dedicated processor or 0.1 processing unit, or 0.05 processing unit when the firmware is at level FW760, or later. Note: HEA is not supported on POWER9 processor-based server.
Memory (physical or logical)	- AIX 5.3 to AIX 6.0: 128 MB - AIX 6.1 or later: 256 MB - Linux: 128 MB	2 GB plus 40 MB for each active logical Host Ethernet Adapter (LHEA)

Table 17. Minimum hardware requirements for logical partitions (continued)

Minimum requirement	AIX and Linux	IBM i
I/O	- Physical or virtual storage adapter (SCSI card) - Physical or virtual network adapter - Storage: - AIX: 2 GB - Linux: Approximately 1 GB	- Load source - Physical or virtual disk I/O adapter (IOA) - Physical or virtual disk unit that is at least 17 GB - Console: your choice of one of the following console types: - Hardware Management Console (HMC) 5250 emulation (requires an HMC) - Operations Console: Requires a LAN connection that supports Operations Console connections. The Operations Console LAN connection can be an embedded port or a LAN IOA. - Alternative restart device: your choice of tape or optical. These devices connect to either a physical or virtual storage adapter. The optical device can be a physical or virtual device. - Physical or virtual LAN adapter that can be used for serviceable event reporting and connection monitoring. At least one IBM i logical partition in the managed system must have a physical LAN adapter that the IBM i logical partition can use for serviceable event reporting and connection monitoring. You can then create a virtual LAN that connects the IBM i logical partition with the physical LAN adapter to the other logical partitions on the managed system, and bridge the physical LAN adapter to the virtual LAN. If the system is managed by using an HMC, the physical LAN adapter must be able to communicate with the HMC so that serviceable event reporting can be routed through the HMC.

Partitioning with the HMC

The *Hardware Management Console (HMC)* is a system that controls managed systems, including the management of logical partitions and use of Capacity Upgrade on Demand. Using service applications, the HMC communicates with managed systems to detect, consolidate, and forward information to IBM for analysis.

The HMC features a browser-based user interface. You can use the HMC locally by connecting a keyboard and mouse to the HMC. You can also configure the HMC so that you can connect to the HMC remotely using a supported browser.

Creating logical partitions

You can create an AIX, Linux, or IBM i logical partition by clicking **Create Partition**, or by using the **Create a Partition from Template** wizard.

About this task

For more information about creating a logical partition by using the **Create a Partition from Template** wizard, see [Creating a logical partition by using a template](#).

To create an AIX, Linux, or IBM i logical partition by using the **Create Partition** option, see [Creating logical partitions by using Create partition](#).

What to do next

When the Hardware Management Console (HMC) is at Version 9.1.0, or later, you can use the **mksyscfg** command to create a logical partition that supports the Physical Page Table (PPT) ratio. During Live Partition Mobility, the Physical Page Table (PPT) ratio is used to translate effective addresses to physical real addresses. PPT is the ratio of the physical memory of the partition and is used by the hypervisor for paging during Live Partition Mobility. To view the PPT ratio attribute of the logical partition, run the **lshwres** command.

Related information

[Changing a partition template to disable Live Partition Mobility](#)

[Viewing system event logs for the Live Partition Mobility disable operation](#)

Creating additional logical partitions

You can use the Create Logical Partition wizard on the Hardware Management Console (HMC) to create a new logical partition. When you create a logical partition, you also create a partition profile that contains the resource allocations and settings for the logical partition.

Before you begin

Use this procedure only if you are creating logical partitions on a managed system that has already been partitioned. If you are creating logical partitions on a new or nonpartitioned managed system, you must test the hardware on your managed system to ensure that the hardware is in working order. Testing the hardware helps you detect potential problems with your managed system and makes such problems easier to correct.

If you plan to create logical partitions that use shared memory, you must first configure the shared memory pool. For instructions, see [“Configuring the shared memory pool” on page 120](#).

If you plan to create AIX logical partitions that use Active Memory Expansion, you must first enable Active Memory Expansion for the server by entering an activation code. For instructions, see [“Entering the activation code for Active Memory Expansion” on page 124](#).

If you want to assign single root I/O virtualization (SR-IOV) logical ports to a logical partition during partition creation, verify whether the managed system supports SR-IOV before you create the logical partition.

About this task

For more information about creating a logical partition and a partition profile on your server by using the HMC, see [“Creating logical partitions” on page 79](#)

What to do next

After creating your logical partition and partition profile, you must install an operating system. For installation instructions for the AIX, IBM i, and Linux operating systems, see [Working with operating systems and software applications for POWER9 processor-based systems](#). For installation instructions for the Virtual I/O Server, see [Installing the Virtual I/O Server and client logical partitions](#).

Related information

[Changing a partition template to disable Live Partition Mobility](#)

[Viewing system event logs for the Live Partition Mobility disable operation](#)

Creating logical partitions on a new or nonpartitioned server

Use these procedures to create logical partitions on your new or nonpartitioned server using the Hardware Management Console (HMC).

When you receive your server, the server is in what is known as the manufacturing default configuration. You can install an operating system on the server and use the server in a nonpartitioned configuration. However, if you want to create logical partitions on the managed system, you must develop a logical partition plan for the server, add hardware to the server or move the hardware within the server according

to your logical partition plan, and validate the hardware on the server. When the server is ready, you can then create the logical partitions using the HMC.

The procedure used to create logical partitions on a new or nonpartitioned server varies by server type.

Assigning a single root I/O virtualization logical port to a logical partition

You can assign a single root I/O virtualization (SR-IOV) logical port to a logical partition by using the Hardware Management Console (HMC).

About this task

For information about managing hardware virtualized I/O adapters, see [Managing hardware virtualized I/O adapters](#).

Creating a logical partition with synchronization of the current configuration

You can create an AIX or LinuxLinux logical partition with synchronization of the current configuration capability by using the Hardware Management Console (HMC).

Prerequisites and assumptions

Ensure that the following prerequisite tasks have been completed before you start the configuration steps:

1. The HMC is set up and configured. For instructions, see [Installing and configuring the HMC](#).
2. You read and understand the [“Logical partition overview”](#) on page 2.
3. You completed the tasks recommended for logical partition planning. For instructions, see [“Planning for logical partitions”](#) on page 61.
4. You removed the system from the manufacturing default configuration and moved the physical hardware to support a partitioned configuration.
5. You have logged in to the HMC with one of the following user roles:
 - Super administrator
 - Operator
6. Ensure that the HMC is at Version 7 Release 7.8.0, or later.

Configuration steps

For more information about creating a logical partition when the HMC is at version 8.7.0, or later, see [“Creating logical partitions”](#) on page 79.

When you create a logical partition by using the **Create Partition** option, or by using the **Create a Partition from Template** wizard, the synchronization of the current configuration is enabled by default.

Enabling the synchronization of the current configuration capability

You can enable the synchronization of the current configuration capability on a logical partition by using the Hardware Management Console (HMC), after the logical partition is created.

Before you begin

Before you plan to enable the feature, ensure that the HMC is at Version 7 Release 7.8.0, is later.

About this task

For information about enabling the synchronization of the current configuration capability on a logical partition after the logical partition is created, see [Changing advanced partition settings](#).

Remote restart states

A simplified remote restartable partition goes through several state changes regarding the simplified remote restart operation, both on the source and destination servers. Most of the simplified remote restart operations are supported only when the partition is in the appropriate remote restart state. A remote restart state is not related to the logical partition state, but is an indicator that is specifically associated with the simplified remote restart operation.

You can use the **lssyscfg** command to view the remote restart status of the partition. The following are the possible values:

Invalid

The partition that is configured for simplified remote restart is in the **Invalid** state until the logical partition is activated. Simplified remote restart is supported only on a logical partition that has been started at least once.

Remote Restartable

After the partition is started and running, the partition transitions into the **Remote Restartable** state. A partition in this state can be remote restarted.

Source Remote Restarting

During the actual simplified remote restart operation, the source partition is in the **Source Remote Restarting** state. This state is transitional and valid until the simplified remote restart operation completes or the operation is canceled.

Destination Remote Restarting

During the actual simplified remote restart operation, the destination partition transitions into the **Destination Remote Restarting** state. This state is transitional and is valid until the simplified remote restart operation completes or the operation is canceled.

Destination Remote Restarted

When the simplified remote restart operation reaches the no return point on the destination system (all the adapters are configured on the destination system), the remote restart status is set to **Destination Remote Restarted**.

Remote Restarted

After the simplified remote restart operation completes, the source logical partition is in the **Remote Restarted** state. The source logical partition can be cleaned up and the destination logical partition is now again ready to be restarted as needed.

Local Storage Update Failed

When an update to persisted information (configuration data stored external to the server on persistent storage) on the Hardware Management Console (HMC) fails due to any reason, the logical partition is in the **Local Storage Update Failed** state. This state indicates that the persisted information on the HMC is out of synchronization with the current logical partition configuration. Simplified remote restart is not allowed in this remote restart state. However, you can use the *usecurrdata* option with the **rrstartlpar** command to run a simplified remote restart operation.

Forced Source Side Restart

When you use the *usecurrdata* option with the **rrstartlpar** command to run a simplified remote restart operation, the partition is restarted with the configuration data on the device and the remote restart state on the source system is updated to **Forced Source Side Restart** state.

Partial Update

When a system is connected to an HMC that has logical partitions with the simplified remote restart capability enabled, the HMC automatically collects the configuration information and the data is stored external to the server on persistent storage. Some configuration information like virtual adapter information, requires a Resource Monitoring and Control (RMC) connection to the Virtual I/O Server (VIOS) partitions. Therefore, the HMC waits until the RMC connection is established to collect such information. When the virtual adapter information is not collected for any reason, the remote restart state is set to **Partial Update**.

Stale Data

When a system is connected to an HMC, the remote restart state is set to **Stale Data** if there is configuration information existing for a logical partition before the state is changed to **Partial Update**. Simplified remote restart is not allowed in this remote restart state. However, you can use the *usecurrdata* option with the **rrstartlpar** command to run a simplified remote restart operation.

Out Of Space

When an update to the persisted information fails because there is insufficient space on storage disk of the HMC to store the configuration information, the remote restart state is updated to **Out Of Space**. You can free up space on storage disk of the HMC and run the **refdev** command to recover from this state.

Profile Restored

When a profile restore operation is performed on a system, during the creation of the simplified remote restart capable partition, the remote restart state is set to **Profile Restored**.

Source Side Cleanup Failed

When a cleanup operation performed on the source system after a successful simplified remote restart fails, remote restart state on the source partition is set to **Source Side Cleanup Failed**.

Reserved Storage Device Update Failed

This state is specific to the simplified remote restart operation that requires a reserved storage device. When an update to the reserved storage device fails for any reason, the logical partition is in the **Reserved Storage Device Update Failed** state. This state indicates that the data on the device is not synchronized with the current partition configuration. Simplified remote restart is not allowed in this remote restart state. However, you can use the *-force* option to run a simplified remote restart operation.

Recovering a simplified remote restart operation

If a simplified remote restart operation of an AIX, Linux, or IBM i logical partition fails, the Hardware Management Console (HMC) attempts an auto recover operation. When the auto recover operation fails, you can recover a simplified remote restart operation by using the HMC command-line interface.

Procedure

To recover a simplified remote restart operation, on the HMC command line, type the following command:

```
rrstartlpar -o recover -m <source server> -t <destination server> -p <lpar name>  
| --id <lpar id> [--force]
```

Aborting a simplified remote restart operation

You can abort or cancel a simplified remote restart operation of an AIX, Linux, or IBM i logical partition by using the Hardware Management Console (HMC) command-line interface.

Procedure

To abort or cancel the simplified remote restart operation of a logical partition, type the following command from the HMC command line:

```
rrstartlpar -m <source managed system> -t <target managed system> -p <lpar name>  
-ip <IP address> [-u <user ID>] -o cancel
```

What to do next

After the cancel operation completes on the destination server, the remote restart status on source server changes to **Remote Restartable**.

Related information

[rrstartlpar command](#)

Viewing the details of a simplified remote restart operation

You can view the details of a simplified remote restart operation of an AIX, Linux, or IBM i logical partition, or the server by using the Hardware Management Console (HMC) command-line interface.

Before you begin

You must ensure that the server is in the Standby or Operating states when you are running the **lsrrstartlpar** command for server-level details. The **lsrrstartlpar** command can be run at the logical partition level for all system states that support the simplified remote restart feature.

Procedure

1. To view the details of the simplified remote restart operation of the server, type the following command from the HMC command line:

```
lsrrstartlpar -r sys -m <managed system>
```

2. To view the details of the simplified remote restart operation of a logical partition, type the following command from the HMC command line:

```
lsrrstartlpar -r lpar -m <managed system> [--filter "lpar_names=" | "lpar_ids=" """]
```

3. To verify whether the auto cleanup feature is enabled, type the following command from the HMC command line:

```
lsrrstartlpar -r -mc <managed system>
```

4. To view the configuration information of all logical partitions that support the simplified remote restart feature, type the following command from the HMC command line:

```
lsrrstartlpar -r lparcfg -m <managed system>
```

Where *lparcfg* shows the logical partition configuration details from the simplified remote restart data.

5. You can view the possible and the recommended mappings of the virtual Fibre Channel adapters, virtual SCSI adapters, and virtual Network Interface Controllers (vNICs) of a logical partition, to Virtual I/O Servers on the destination server, by typing the following command from the HMC command line:

```
lsrrstartlpar -r virtualio -o validate -m <source-system-name> -t <target-system-name>  
[--vniccfg <1|2> ] [-- redundantvnicbkdev <1|2>]  
[--ip <IP address>] [-u <user ID> ]  
--filter "<filter data>"
```

Where

- *virtualio* lists all the available Virtual I/O adapters on the destination system that can be used for mapping virtual I/O.
- *vniccfg* specifies whether the vNIC configuration must be maintained. A value of 1 indicates that the vNIC configuration must be maintained. If you specify the value as 1, and if the vNIC configuration cannot be maintained, the remote restart operation fails. A value of 2 indicates that the vNIC configuration must be maintained when possible. If you specify the value as 2, and if the vNIC configuration cannot be maintained, the remote restart operation succeeds but the vNIC configuration of the logical partition is not retained after the remote restart operation completes.
- *redundantvnicbkdev* specifies whether the vNIC backing device redundancy must be maintained. A value of 1 indicates that the vNIC backing device redundancy must be maintained. If you specify the value as 1, and if the vNIC backing device redundancy cannot be maintained, the remote restart operation fails. A value of 2 indicates that the vNIC backing device redundancy must be maintained when possible. If you specify the value as 2, and if the vNIC backing device redundancy cannot be maintained, the remote restart operation succeeds but the vNIC backing device redundancy of the logical partition is not retained after the remote restart operation completes.
- *ip* is the IP address or the host name of the management console of the destination server.

- *u* is the user ID that must be used on the management console of the destination server.
- *filter* is a required parameter and only one logical partition (either partition ID or partition name) can be specified when the *virtualio* parameter is specified.

Related information

[lsrrstartlpar command](#)

Validating the simplified remote restart operation of a logical partition

You can validate the simplified remote restart operation of an AIX, Linux, or IBM i logical partition by using the Hardware Management Console (HMC) command-line interface.

Before you begin

When the source and destination servers are managed by different Hardware Management Consoles, you must ensure that the HMC at the source server and the destination server is at version 8.5.0 or later.

Note: When the HMC is at Version 9.2.950, or later, and the firmware is at a level FW950, or later, for the remote restart operation to be successful on the logical partition that supports platform keystore capability, the user-defined system key that is configured on both the source and the destination system must match.

Procedure

1. To validate the simplified remote restart operation of a logical partition, type the following command from the HMC command line:

```
lsrrstartlpar -m <source managed system> -t <target managed system> -p <lpar name> |
--id <lpar id> -o validate
```

To validate the simplified remote restart operation of a logical partition when the source server and the destination server are managed by a different HMC, type the following command:

```
lsrrstartlpar -m <source managed system> -t <target managed system> -p <lpar name> |
--id <lpar id> --ip <IP address> [-u <user ID>] -o validate
```

Where

- *IP address* is the IP address or the host name of the HMC that manages the destination server.
- *user ID* is the user ID that is used on the HMC that manages the destination server.

If the HMC at the destination server is not at version 8.5.0 or later, the validation operation fails and error messages are displayed for your appropriate action.

2. To specify the destination shared processor pool, type the following command from the HMC command line:

```
lsrrstartlpar -m <source managed system> -t <target managed system> -p <lpar name>
| --id <lpar id> --ip <IP address> [-u <user ID>] - i|f
"shared_proc_pool_name=<spp name>|shared_proc_pool_id=<spp id>" -o validate
```

Where

- *shared_proc_pool_name* is the name of the shared processor pool on the destination server.
- *shared_proc_pool_id* is the ID of the shared processor pool.

Note: The *shared_proc_pool_id* and *shared_proc_pool_name* attributes are mutually exclusive.

3. To specify the destination virtual Fibre Channel mapping, type the following command from the HMC command line:

```
lsrrstartlpar -m <source managed system> -t <target managed system> -p <lpar name>
| --id <lpar id> --ip <IP address> [-u <user ID>] - i|f "virtual_fc_mappings
=slot_num/vios_lpar_name/vios_lpar_id/[vios_slot_num]/[vios_fc_port_name]" -o validate
```

Where

- *slot num* is virtual Fibre Channel slot number.
 - *vios_slot_num* is virtual Fibre Channel slot number of the VIOS.
4. To validate the simplified remote restart operation of a logical partition with a different number of dedicated processors or virtual processors on the destination server than what the partition was assigned on the source server, or with a different number of shared processing units on the destination server than what the logical partition was assigned on the source server, type the following command from the HMC command line:

```
rrstartlpar -o validate -m <managed system> -t <target system> [-p <partition name>
| --id <partition ID>] -i "desired_procs = <desired procs absolute value>|min,
desired_proc_units = <desired proc units absolute value>|min"
```

The value for the *desired_procs* attribute indicates the number of processors with which logical partition can be restarted. The following are the possible values for this attribute:

- *desired procs absolute value* - Must be greater than or equal to the current minimum dedicated processors or virtual processors, and less than or equal to current maximum dedicated processors or virtual processors.
 - *desired proc units absolute value* - Must be greater than or equal to the current minimum processing units, and less than or equal to current maximum processing units.
 - *min* - The value of the current minimum processing units, or the current minimum dedicated processors or virtual processors. In the case of dedicated processors, the logical partition is started with the minimum number of dedicated processors or virtual processors the logical partition was assigned on the source server. In the case of shared processors, the logical partition is started with the minimum number of shared processing units the partition was assigned on the source server.
5. To validate the simplified remote restart operation of a logical partition when you want to start the logical partition with a different amount of memory on the destination server than what the logical partition was assigned on the source server, type the following command from the HMC command line:

```
rrstartlpar -o validate -m <managed system> -t <target system> [-p <partition name>
| --id <partition ID>] -i "desired_mem = <desired absolute memory>|min"
```

The value for the *desired_mem* attribute indicates the amount of memory with which the logical partition can be restarted. The following are the possible values for this attribute:

- *desired_absolute_memory* - Must be greater than or equal to the current minimum memory, and less than or equal to current maximum memory.
 - *min* - The value of the current minimum memory. The logical partition is started with the minimum amount of memory the logical partition was assigned on the source server.
6. To validate the simplified remote restart operation of a logical partition when you want to start the logical partition with a different virtual switch on the destination server than what the logical partition was assigned on the source server, type the following command from the HMC command line:

```
rrstartlpar -o validate -m <managed system> -t <target system> [-p <partition name>
| --id <partition ID>] -i "vswitch_mappings = <vlanId_1>/<source_vswitchName_1>/
<target_vswitchName_1>, ..., <vlanId_n>/<source_vswitchName_n>/<target_vswitchName_n>"]"
```

Related information

[rrstartlpar command](#)

Remotely restarting a logical partition

You can restart an AIX, Linux, or IBM i logical partition remotely by using the Hardware Management Console (HMC) command-line interface. You can run up to four concurrent simplified remote restart operations for a destination server. When the HMC is at version 8.5.0 or later, you can run up to 32 concurrent simplified remote restart operations for a destination server.

Before you begin

Before you begin, complete the following tasks:

- You must ensure that the HMC is at version 8.2.0 or later, and the firmware is at level FW820, or later, and the VIOS must be at version 2.2.3.4 with the VIOS interim fix IV63331m4a or later, for the simplified remote restart feature.
- You must ensure that the source and the destination servers are connected to the same HMC. However, when the HMC is at version 8.5.0 or later, the source and destination server can be connected to a different Hardware Management Consoles. In this case, the HMC at the source and destination servers must be at version 8.5.0, or later.
- When the HMC is at version 8.3.0, or later, you can restart a logical partition on another server only when the source server is in the Initializing, Power Off, Error, or Error - Dump in progress state and the destination server is in the Operating state. When the HMC is at version 8.4.0, or later, you can restart a logical partition on another server only when the source server is in the Initializing, Power Off, Powering Off, Error, or Error - Dump in progress in progress state and the destination server is in the Operating state. When the HMC is at version 8.5.0, or later, you can restart a logical partition on another server only when the source server is in the Initializing, Power Off, Powering Off, No Connection, Error, or Error - Dump in progress state, and the destination server is in the Operating state.
- When the HMC is at Version 9.2.950, or later, and the firmware is at a level FW950, or later, for the remote restart operation to be successful on the logical partition that supports platform keystore capability, the user-defined system key that is configured on both the source and the destination system must match.

Procedure

1. To remotely restart a logical partition, type the following command from the HMC command line:

```
rrstartlpar -o restart -m <source managed system> -t <target managed system> -p <lpar name>
```

When the source server and the destination server are managed by different Hardware Management Consoles, type the following command to remotely restart a logical partition:

```
rrstartlpar -m <source managed system> -t <target managed system> -p <lpar name> | --id  
<lpar id>  
--ip <IP address> [-u <user ID>] -o restart
```

Note: If a logical partition that is enabled with the simplified version of the remote restart feature was in the Suspended state before the simplified remote restart operation started, the simplified remote restart operation fails. You can use the `--force` option to force a simplified remote restart operation.

To verify the status of the simplified remote restart operation, type the following command from the HMC command line:

```
lssyscfg -r lpar -m <server> -F name,state,remote_restart_status
```

You can also run the **lsrrstartlpar** command to view the status of the simplified remote restart operation.

2. To specify the destination shared processor pool for the simplified remote restart operation, type the following command from the HMC command line:

```
rrstartlpar -m <source managed system> -t <target managed system> -p <lpar name>  
| --id <lpar id> --ip <IP address> [-u <user ID>] -i |f  
"shared_proc_pool_name=<spp name>|shared_proc_pool_id=<spp id>" -o restart
```

Where:

- *shared_proc_pool_name* is the name of the shared processor pool on the destination server.
- *shared_proc_pool_id* is the ID of the shared processor pool.

Note: The *shared_proc_pool_id* and *shared_proc_pool_name* attributes are mutually exclusive.

3. To specify the destination virtual Fibre Channel mapping for the simplified remote restart operation, type the following command from the HMC command line:

```
rrstartlpar -m <source managed system> -t <target managed system> -p <lpar name> | --id  
<lpar id>  
--ip <IP address> [-u <user ID>] -i |f "virtual_fc_mappings=  
slot_num/vios_lpar_name/vios_lpar_id/[vios_slot_num]/[vios_fc_port_name]" -o restart
```

Where:

- *slot num* is virtual Fibre Channel slot number.
 - *vios_slot_num* is virtual Fibre Channel slot number of the VIOS.
4. To remotely restart a partition when the source server is in the No Connection state, type the following command from the HMC command line:

```
rrstartlpar -m <source managed system> -t <target managed system> -p <lpar name> | --id  
<lpar id> -o restart --noconnection
```

5. When the HMC is at version 8.5.0, or later, an auto cleanup operation on the server is performed by the HMC after the successful completion of the simplified remote restart operation. The source server is back to the Standby or Operating state, and the Resource Monitoring and Control (RMC) connection for the Virtual I/O Server (VIOS) partitions that are serving the remotely restarted logical partitions becomes active.

You can enable the auto cleanup operation by typing the following command from the HMC command line:

```
rrstartlpar -o set -r <source managed system> -i "auto_cleanup_enabled=1"
```

You can disable the auto cleanup operation by typing the following command from the HMC command line:

```
rrstartlpar -o set -r <source managed system> -i "auto_cleanup_enabled=0"
```

The auto cleanup operation is enabled by default. If the remote restarted partitions are not cleaned up automatically, type the following command from the HMC command line to perform a cleanup operation on the source server:

```
rrstartlpar -o cleanup -m source managed system -p lpar name
```

The default behavior of a cleanup operation is to remove the reserved storage device from the device pool on the source server. You can use the *--retaindev* option when you want to retain the reserved storage device in the pool and override the default action of the cleanup operation.

6. To remotely restart a logical partition with a different number of dedicated processors or virtual processors on the destination server than what the partition was assigned on the source server, or with a different number of shared processing units on the destination server than what the logical partition was assigned on the source server, type the following command from the HMC command line:

```
rrstartlpar -o restart -m <managed system> -t <target system> [-p <partition name>  
| --id <partition ID>] -i "desired_procs = <desired procs absolute value>|min,  
desired_proc_units =  
<desired proc units absolute value>|min"
```

The value for the *desired_procs* attribute indicates the number of processors with which the logical partition can be restarted. The following are the possible values for this attribute:

- *desired_procs absolute value* - Must be greater than or equal to the current minimum dedicated processors or virtual processors, and less than or equal to current maximum dedicated processors or virtual processors.
- *desired_proc units absolute value* - Must be greater than or equal to the current minimum processing units, and less than or equal to current maximum processing units.

- *min* - The value of the current minimum processing units, or the current minimum dedicated processors or virtual processors. In the case of dedicated processors, the logical partition is started with the minimum number of dedicated processors or virtual processors the logical partition was assigned on the source server. In the case of shared processors, the logical partition is started with the minimum number of shared processing units the partition was assigned on the source server.
7. To remotely restart a logical partition when you want to start the logical partition with a different amount of memory on the destination server than what the logical partition was assigned on the source server, type the following command from the HMC command line:

```
rrstartlpar -o restart -m <managed system> -t <target system> [-p <partition name> |
--id <partition ID>] -i "desired_mem = <desired absolute memory>|min"
```

The value for the *desired_mem* attribute indicates the amount of memory with which the logical partition can be restarted. The following are the possible values for this attribute:

- *desired_absolute_memory* - Must be greater than or equal to the current minimum memory, and less than or equal to current maximum memory.
 - *min* - The value of the current minimum memory. The logical partition is started with the minimum amount of memory the logical partition was assigned on the source server.
8. To remotely restart a logical partition when you want to start the logical partition with a different virtual switch on the destination server than what the logical partition was assigned on the source server, type the following command from the HMC command line:

```
rrstartlpar -o restart -m <managed system> -t <target system> [-p <partition name> |
--id <partition ID>] -i "vswitch_mappings = <vlanId_1>/<source_vswitchName_1>/
<target_vswitchName_1>, ..., <vlanId_n>/<source_vswitchName_n>/<target_vswitchName_n>"]"
```

9. When you want to prevent a logical partition from being started during the simplified remote restart operation, type the following command from the HMC command line:

```
rrstartlpar -o restart -m <managed system> -t <target system> [-p <partition name>
| --id <partition ID>] --skippoweron
```

10. You can test the simplified remote restart operation of a shutdown partition from the source server, which is in Operating or Standby state by typing the following command from the HMC command line:

```
rrstartlpar -o restart -m <managed system> -t <target system> [-p <partition name>
| --id <partition ID>] --test
```

Related information

[rrstartlpar command](#)

Enabling or disabling the simplified remote restart capability

You can enable or disable the simplified remote restart capability of a logical partition after the logical partition is created, by using the Hardware Management Console (HMC).

Before you begin

- Ensure that the HMC is at Version 8.1.0, or later.
- Ensure that the server is a POWER9 processor-based server that supports partitions that are capable of simplified remote restart.
- Ensure that the logical partition is in the Not Activated state. When the HMC is at Version 8.6.0, or later, and the firmware is at level FW860, or later, you can enable or disable the simplified remote restart capability when the logical partition is in the Running state. The logical partition must not be in the Suspended, Resuming, Migrating, or Remote Restarting states.
- To use the simplified remote restart feature, ensure that the HMC is at Version 8.2.0, or later.

About this task

To enable or disable the simplified remote restart capability of a logical partition by using the HMC, complete the following steps:

Procedure

1. From the HMC command line, type the following command to enable the simplified remote restart feature:

```
chsyscfg -r lpar -m managed-system -i "name=partition name, remote_restart_capable=1"
```

From the HMC command line, type the following command to enable the simplified remote restart feature:

```
chsyscfg -r lpar -m managed-system -i "name=partition name,  
simplified_remote_restart_capable=1"
```

2. From the HMC command line, type the following command to disable the simplified remote restart feature:

```
chsyscfg -r lpar -m managed-system -i "name=partition name, remote_restart_capable=0"
```

From the HMC command line, type the following command to disable the simplified remote restart feature:

```
chsyscfg -r lpar -m managed-system -i "name=partition name,  
simplified_remote_restart_capable=0"
```

Creating a logical partition with the simplified remote restart capability

You can create an AIX, IBM i, or Linux logical partition with the simplified remote restart capability by using the Hardware Management Console (HMC). The HMC provides options to enable the simplified remote restart of the logical partition when the logical partition is created.

Prerequisites and assumptions

Ensure that the following prerequisite steps have been completed before you start the configuration steps:

- Before you create a logical partition with the simplified remote restart capability, verify the following requirements:
 1. You read the configuration requirements and restrictions for creating a logical partition with the simplified remote restart capability. For instructions, see [“Configuration requirements and restrictions for the remote restart capability of a logical partition”](#) on page 67.
 2. Before you create a logical partition with the simplified remote restart capability, verify the following requirements:
 - Verify that the server supports the simplified remote restart capability. For instructions, see [“Verifying that the server supports partitions that are capable of the simplified version of the remote restart feature”](#) on page 68.
 - You read the configuration requirements and restrictions for creating a logical partition with the simplified remote restart capability. For instructions, see [“Configuration requirements and restrictions for the remote restart capability of a logical partition”](#) on page 67.

You can create a partition with the simplified remote restart capability by using a template that supports the capability when you create a partition by using the **Create a Partition from Template** wizard. When the template does not support the capability, you can enable the capability by modifying the template and then using the modified template to create a partition by using the **Create a Partition from Template** wizard. For more information about editing a partition template, see [Changing a partition template](#).

For more information about enabling the simplified remote restart feature after you create the logical partition, see [Changing partition properties and capabilities](#).

Creating a logical partition with Virtual Trusted Platform capability

You can create an AIX logical partition with Virtual Trusted Platform Module (VTPM) capability by using the Hardware Management Console (HMC). HMC Version 7.7.4 or later, provides an option to enable a VTPM on the logical partition when the logical partition is created. The HMC also provides an option to enable a VTPM on a running logical partition.

Prerequisites and assumptions

Ensure that the following prerequisite tasks have been completed and are operational before you start the configuration steps:

1. The HMC is set up and configured. For instructions, see [Installing and configuring the HMC](#).
2. You read and understand the “[Logical partition overview](#)” on page 2.
3. You completed the tasks recommended for logical partition planning. For instructions, see “[Planning for logical partitions](#)” on page 61.
4. You removed the system from the manufacturing default configuration and moved the physical hardware to support a partitioned configuration. For instructions, see “[Creating logical partitions on a new or nonpartitioned server](#)” on page 80.
5. You verified that the server has logical partition support for VTPM. For instructions, see “[Verifying that the server supports Virtual Trusted Platform Module](#)” on page 69

You can also enable the VTPM on a logical partition after logical partition creation. For instructions, see “[Enabling and disabling a Virtual Trusted Platform Module on a logical partition](#)” on page 91

6. You have logged in to the HMC with one of the following user roles:

- Super administrator
- Operator

You can create a partition with the VTPM capability by using a template that supports the capability when you create a partition by using the **Create a Partition from Template** wizard. When the template does not support the capability, you can enable the capability by modifying the template and then using the modified template to create a partition by using the **Create a Partition from Template** wizard. For more information about editing a partition template, see [Changing a partition template](#). For more information about enabling the VTPM capability after partition creation, see [Changing partition properties and capabilities](#).

Enabling and disabling a Virtual Trusted Platform Module on a logical partition

You can enable a Virtual Trusted Platform Module (VTPM) on a logical partition by using the Hardware Management Console (HMC), after the logical partition is created.

Before you begin

To enable a VTPM, ensure that an AIX, Linux or a Virtual I/O Server (VIOS) logical partition is in the Not activated state.

About this task

For information about enabling VTPM on a logical partition, see [Changing advanced partition settings](#).

Results

If you dynamically enable a VTPM on a logical partition, the VTPM function is activated only at the next logical partition activation. However, disabling a VTPM takes effect immediately.

What to do next

To dynamically disable a VTPM, log on to the AIX, Linux or VIOS logical partition and disable the Trusted Computing Services daemon (tcsd) by using the **stopsrc** command. When the **tcsd** software is stopped, the device must be removed from the AIX logical partition by using the **rmdev** command. After the device is successfully deleted from the AIX logical partition, use the HMC to clear the VTPM check box from the properties of the partition. This completely removes the device and deletes all stored data that is associated with the VTPM.

Viewing the Virtual Trusted Platform Module settings

You can view the advanced Virtual Trusted Platform Module (VTPM) settings by using the Hardware Management Console (HMC).

About this task

For more information about viewing the VTPM settings, see [Changing advanced partition settings](#).

Creating additional partition profiles

You can create more than one partition profile for a logical partition using the Hardware Management Console (HMC). Each partition profile can specify a different amount of system resources and different logical partition startup attributes. You can change the attributes used by a logical partition by shutting down the logical partition and restarting the logical partition using a different partition profile.

Before you begin

If you plan to create a partition profile in which you configure Active Memory Expansion for an AIX logical partition, ensure that you enter an activation code to enable Active Memory Expansion on the server before you activate the logical partition with this partition profile. For instructions, see [“Entering the activation code for Active Memory Expansion”](#) on page 124.

When you create a partition profile, do not select **Use all the resources in the system** when both of the following conditions are true:

- You plan to create a partition profile that uses all of the system resources.
- You plan to configure Active Memory Expansion for that partition profile.

Instead, manually assign all of the resources in the system to the partition profile. In the process, you can configure the Active Memory Expansion factor.

About this task

To create a partition profile using the HMC, follow these steps:

Procedure



1. In the navigation pane, click the **Resources** icon.
2. Click **All Partitions**. The All Partitions page is displayed.
3. Select the logical partition and click **Actions > Profiles > Manage Profiles**.
4. Click **Actions > New**.
5. Follow the steps in the Create Partition Profile wizard to create the partition profile.

What to do next

If you created at least one virtual fibre channel adapter, complete the following tasks to connect the logical partition to its storage:

1. Activate the logical partition. When you activate the logical partition, the HMC assigns a pair of worldwide port names (WWPNs) to the virtual fibre channel adapter. For instructions, see [“Activating a logical partition”](#) on page 133.
2. Restart the Virtual I/O Server (that provides the connection to a physical fibre channel adapter) or run the **syscfg** command. This enables the Virtual I/O Server to recognize the WWPNs of the virtual fibre channel adapter on the client logical partition. For instructions, see [“Restarting Virtual I/O Server logical partitions by using the HMC”](#) on page 145.
3. Assign the virtual fibre channel adapter on the client logical partition to a physical port of a physical fibre channel adapter. For instructions, see [“Changing virtual Fibre Channel for a Virtual I/O Server by using the HMC”](#) on page 173.

Creating a system profile

You can create a system profile using the Hardware Management Console (HMC). A *system profile* is an ordered list of partition profiles. When you activate a system profile, the managed system attempts to activate the partition profiles in the system profile in the order in which the partition profiles are listed.

Before you begin

System profiles are also useful for validating your partition profiles to ensure that you have not overcommitted the resources on your managed system.

Restriction: You cannot create system profiles that contain logical partitions that use shared memory.

About this task

To create a system profile using the HMC, follow these steps:

Procedure



1. In the navigation pane, click the **Resources** icon.
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **Actions** > **View System Properties**. The **Properties** page is displayed.
4. Select the system and click **System Actions** > **Legacy** > **Manage System Profiles**.
5. Click **Actions** > **New**.
6. Enter the name of the new system profile into **System profile name**.
7. For each partition profile that you want to add to the system profile, open the logical partition to which the partition profile belongs, select the partition profile, and click **Add**.
8. Click **OK**.

Creating an AIX logical partition that uses IBM i virtual I/O resources

You can create an AIX logical partition that uses IBM i virtual I/O resources on servers that are managed by a Hardware Management Console (HMC). This allows you to maximize utilization of the physical hardware and simplify the backup procedure for your managed system.

Before you begin

To set this up, you must create virtual SCSI adapters that connect the AIX logical partition with the IBM i. You can then set up IBM i to provide disk resources to the AIX logical partition through the virtual SCSI connection. You can also create a virtual serial connection between the IBM i logical partition and the AIX logical partition. A virtual serial connection allows you to connect to the AIX logical partition from the IBM i logical partition.

Alternatively, you can create a Virtual I/O Server logical partition and configure the AIX logical partition to use the virtual SCSI and virtual Ethernet resources of the Virtual I/O Server logical partition. You might need to enter a PowerVM Editions PowerVM for IBM PowerLinux activation code to create a Virtual I/O Server logical partition on your server.

About this task

For more information about creating a logical partition, see [“Creating logical partitions”](#) on page 79. For more information about adding IBM i hosted virtual SCSI adapters when the HMC is at version 8.7.0, or later, see [Managing virtual storage for a partition](#).

Creating a network-server description and a network-server storage space for an AIX logical partition

A *network-server description* (NWSD) is an IBM i object that describes the storage resources that are used by an integrated operating environment. An NWSD can be linked to one or more network-server storage spaces. Create an NWSD to assign storage to an AIX logical partition that uses IBM i resources.

About this task

To create an NWSD and a network-server storage space for an AIX logical partition that uses IBM i resources, follow these steps:

Procedure

1. Determine the correct SCSI server resource name.
 - If there is only one SCSI server adapter corresponding to a given client logical partition, and that adapter has its remote logical partition and remote slot configured correctly, you can specify *AUTO as the RSRCTYPE in your NWSD.
 - Otherwise, you must determine the actual resource name. At an IBM i command line, type `WRKHDWRSC *CMN`, and find a controller resource with type 290B and a converged location code that corresponds to the SCSI server adapter at the Hardware Management Console (HMC). This resource name will be used later to specify the SCSI server resource.
2. At an IBM i command line on the logical partition that shares resources, type `CRTNWSD` and press F4 for prompts.
3. Specify the following information:

The default or suggested parameter values are provided within the parentheses. These settings are relevant only to a logical partition. After the installation, if your root file system (/) is not installed on the first partition of the first disk, you must set a root parameter.

- NWSD (Provide a name for the NWSD)
- RSRCTYPE (*AUTO or the resource name of the SCSI server resource)
- TYPE (*GUEST)
- ONLINE (*NO or *YES)
- PARTITION ('Provide the name of your AIX logical partition')

As an alternative to the Partition parameter, you can also specify a logical partition number by typing `PTNNBR(integer)` where *integer* is the number of the logical partition you are specifying.

- CODEPAGE (437)
- TCPPTCFG (*NONE)
- RSTDDEVRSC (for virtual CD and tape devices) (*NONE)
- SYNCTIME (*TYPE)
- IPLSRC (*NWSSTG)

- You can store a kernel in a disk partition of a virtual disk (a network-server storage space (NWSSTG)). By specifying the IPLSRC (*NWSSTG) parameter, you are specifying that the AIX logical partition will start from a disk partition on that virtual disk. The disk partition on the virtual disk must be formatted as type PReP Boot (type 0x41) and marked as a device that starts.
 - To start an NWSD with a kernel from a stream file, set the IPLSRC parameter to *STMF and set the IPLSTMF parameter to point to the kernel. You must have read access to the file and the path leading to the file to use the vary on command. This value only loads the kernel. After the kernel is running, it must find a root file system. In an initial installation, the root file system might be a RAM disk that is physically attached to the kernel.
 - IPLSTMF (*NONE)
 - IPLPARM (*NONE)
 - PWRCTL (*YES)
 - If you specify PWRCTL (*YES), perform the following steps:
 - a. Ensure that the server adapter in the IBM i logical partition specifies the remote logical partition and remote slot in its configuration.
 - b. Ensure that the client logical partition has the IBM i logical partition as the power-controlling logical partition in the profile.
 - c. Ensure before you activate the NWSD that the client logical partition's profile has been saved to the server by activating the logical partition from the HMC, even if the client operating system does not activate correctly because of the absence of virtual devices.
 - If you specify PWRCTL (*NO), virtual devices will be available to the logical partition. You must shut down and restart the logical partition using the HMC.
4. If you use IBM Navigator for i, create the network-server storage space using IBM Navigator for i.
 - a) Expand **My Connections > your server > Network > Windows Administration**.
 - b) Right-click the **Disk Drives** and select **New Disk**.
 - c) In the **Disk drive name** field, specify the name that you want to give to the disk drive.
 - d) In the **Description** field, specify a meaningful description for the disk drive.
 - e) In the **Capacity** field, specify the size of the new disk drive in megabytes.
 - f) Click **OK**.
 - g) Continue with step “6” on page 95.
 5. If you use a character-based interface, create the network-server storage space using a character-based interface:
 - a) At an IBM i command line, type the command CRTNWSSTG and press F4.
The Create NWS Storage Space (CRTNWSSTG) display opens.
 - b) In the Network-server storage space field, specify the name you want to give to the storage space.
 - c) In the Size field, specify the size in megabytes for the new storage space.
 - d) In the Text description field, specify a meaningful description for the storage space.
 - e) Press Enter.
 - f) Continue with step “7” on page 96.
 6. If you use IBM Navigator for i, link the network-server storage space using IBM Navigator for i.
 - a) Expand **My Connections > your server > Network > Windows Administration**.
 - b) Click **Disk Drives**, right-click an available network-server storage space, and select **Add Link**.
 - c) Select the server to which you want to link the network-server storage space.
 - d) Select the link sequence position you want to use.
 - e) Select one of the available data access types.
 - f) Click **OK**.
- The procedure is complete. Do not complete step “7” on page 96.

7. If you use a character-based interface, link the network-server storage space using a character-based interface:
 - a) At an IBM i command line, type the command ADDNWSSTGL and press F4.
The Add Network-Server Storage Link (ADDNWSSTGL) display opens.
 - b) In the Network server description field, specify the name of the network server description (NWSD).
 - c) In the Dynamic storage link field, specify *YES to make the network-server storage space dynamically available to the logical partition (that is, available without rebooting the AIX logical partition).
 - d) In the Drive sequence number field, specify the link sequence position you want to use.
 - e) Press Enter.

Connecting an AIX logical partition to the virtual console

You can connect an AIX logical partition to the virtual console so that you can install the operating system or access the command line interface for the AIX logical partition.

Before you begin

You must have one of the following privileges to use the AIX virtual console:

- System Partitions - Operations
- System Partitions - Administration

About this task

The virtual console provides console functions for a AIX server. Use it primarily during initial operating system installation. You can also use the virtual console to view server errors or restore local area network (LAN) communication. Use this console connection before you configure the Transmission Control Protocol (TCP) or Internet Protocol (IP).

You can use any Telnet client as the AIX console. Multiple Telnet clients can share access to the same virtual console. TCP/IP must be configured and running on at least one IBM i logical partition.

For IBM i version 7.4, IBM i version 7.5 Technology Refresh 8, and IBM i version 7.6 Technology Refresh 2, the virtual console server supports the following operational modes:

- Unsecured mode: The server accepts direct Telnet connections on port 2301 from any network interface. The unsecured mode is the default setting for IBM i version 7.4, IBM i version 7.5, and IBM i version 7.6.
- Secure mode: The server accepts connections only from the localhost 127.0.0.1:2301 port. To access the server remotely, create an SSH tunnel with the local port forwarding capability, and then connect with a Telnet client to the forwarded local port. SSH encrypts all traffic, including login credentials.

Default security modes by IBM i release:

- Version 7.4: Unsecured mode
- Version 7.5 Technology Refresh 8: Unsecured mode
- Version 7.6 Technology Refresh 2: Unsecured mode

Using secure mode to connect to the virtual console

In the secure mode, create an SSH tunnel before you connect to the virtual console. The SSH tunnel encrypts all traffic, including service tools credentials.

Prerequisites and privileges for secure mode:

- OpenSSH is installed and the OpenSSH server is running on IBM i. Install OpenSSH with license program 5733-SC1 Option 1.
- PASE must be installed. Install PASE with 5770-SS1 Option 33.

- The SSH client is available on the workstation and installed in AIX, Linux, Mac, and Windows version 10 or 11
- IBM i users must have the Secure Shell (SSH) access
- Service tool user ID must have System partitions - operations or System partitions - administration access.
- For the IBM i version 7.4, the PTF MJ09039 must be installed for security mode changes
- For IBM i version 7.5 Technology Refresh 8, the PTF MJ08967 must be installed for security mode changes
- For IBM i version 7.6 Technology Refresh 2, the PTF MJ08968 must be installed for security mode changes

To connect to a virtual console through the secure mode, complete the following steps:

1. Establish the SSH tunnel:

- a. Open the console or the command line on your workstation and use the local port forwarding process to create an SSH tunnel. An example command using a profile with the SSH access follows:

```
ssh -L localhost:23011:localhost:2301 IBMiUserWithSSHAccess@IBMiHostname
```

- b. Enter the password of the IBM i profile in step 1a.

Note: Do not close the SSH session.

- c. Follow step 2 to connect with a standard Telnet Client, or step 3 to connect with the IBM Personal Communications client.

2. Connect with a standard Telnet client:

- a. Open a second console or command line on your workstation, separate from the SSH tunnel, and connect your Telnet client to the local forwarded port you selected in Step 1a. An example command follows:

```
telnet localhost <your_chosen_port>
```

Note: In the example in step 1a, the **your_chosen_port** value is 23011.

- b. Select the logical partition for which you want to connect to the virtual console.

- c. Enter the IBM i service tools ID and password.

Note: If Multi-Factor Authentication (MFA) is configured for the service tools ID, enter the One-Time Password (TOTP) when prompted.

- d. The telnet client is now securely connected to the console of the selected logical partition. After you finish using the console, close the Telnet session and then enter **exit** in the SSH console to close the SSH connection. This completes the procedure of connecting to a virtual console through the secure mode.

3. Connect with the IBM Personal Communications:

- a. Click **Start > IBM Personal Communications > Start or Configure Session**. The **Customize Communication** page is displayed.
- b. Select **ASCII** as your type of host and select **Link Parameters**. The **Telnet ASCII** page is displayed.
- c. Enter the hostname or the IP address of the logical partition that is sharing its resources, and enter the local port number that you chose in step 1a.

Note: In the example in step 1a, the **your_chosen_port** value is 23011.

- d. Click **OK**.

- e. From the **Guest Partition Consoles** page, select the logical partition to which you want to connect as the console.

- f. Enter the IBM i service tools ID and password to connect to the AIX logical partition.

Note: If Multi-Factor Authentication (MFA) is configured for the service tools ID, enter the One-Time Password (TOTP) when prompted.

- g. The IBM Personal Communications telnet client is now securely connected to the console of the selected logical partition. After you finish using the console, close the Telnet session and then enter `exit` in the SSH console to close the SSH connection. This completes the procedure of connecting to a virtual console through the secure mode.

Note: When you run the local x5250 virtual console session with the **TELNET** Command Language (CL) command to connect to the local host system, the console displays all typed characters on the screen, including service tools credentials. This behavior creates a security risk during screen share sessions where other users can view your credentials. To mitigate this risk, use a separate SSH tunnel with a standard telnet client or IBM Personal Communications telnet client, even when you connect from the local host system.

Using unsecured mode to connect to the virtual console

The unsecured mode transmits IBM i credentials and console traffic in an unencrypted form across the network. Use this mode only on local networks where security is not a concern.

Note: In IBM i versions before 7.4 with PTF MJ09039, 7.5 Technology Refresh 8, and 7.6 Technology Refresh 2 versions, the console server accepts telnet connections from all network interfaces, including the localhost. You can still use the SSH tunnel method to establish a secure connection by encrypting the telnet login credentials and the connection data on the network. For more information see [Using secure mode to connect to the virtual console](#).

To connect to a virtual console through the unsecured mode, complete the following steps:

Note: Follow step 1 to connect with a standard Telnet Client, or step 2 to connect with IBM Personal Communications.

1. Connect with a standard Telnet client:

- a. Open a console or command prompt on your workstation.
- b. Run the **Telnet** command to connect your IBM i system using port 2301. Example:

```
telnet IBMiHostname 2301
```

Note: The `IBMiHostname` is the IBM i system hostname or IP address.

- c. Select the logical partition for which you want to connect to the virtual console.
- d. Enter the IBM i service tools ID and password.

Note: If Multi-Factor Authentication (MFA) is configured for the service tools ID, enter the One-Time Password (TOTP) when prompted.

- e. The telnet client is now connected to the console of the selected logical partition. After you finish using the console, close the Telnet session. This completes the procedure of connecting to a virtual console through the unsecure mode.

2. Connect with the IBM Personal Communications:

- a. Click **Start > IBM Personal Communications > Start or Configure Session**. The **Customize Communication** page is displayed.
- b. Select **ASCII** as your type of host and select **Link Parameters**. The **Telnet ASCII** page is displayed.
- c. Enter the hostname or the IP address of the logical partition that is sharing its resources, and enter port number 2301.
- d. Click **OK**.
- e. From the **Guest Partition Consoles** page, select the logical partition to which you want to connect as the console.
- f. Enter the IBM i service tools ID and password.

Note: If Multi-Factor Authentication (MFA) is configured for the service tools ID, enter the One-Time Password (TOTP) when prompted.

- g. The IBM Personal Communications telnet client is now connected to the console of the selected logical partition. After you finish using the console, close the Telnet session. This completes the procedure of connecting to a virtual console through the unsecure mode.

Managing the virtual console security mode

Use the **VIOFR** command to manage virtual console security mode and switch between secure and unsecured modes without an Initial Program Load (IPL). Run the **VIOFR** command from System Service Tools (SST) or Dedicated Service Tools (DST).

You must have the following privileges to run the **VIOFR** command:

- Service tools user ID
- Display/Alter/Dump privilege
- Start Service tool privilege
- Service tool security privilege

To run the **VIOFR** command from SST, complete the following steps:

1. In the 5250 command session, enter STRSST and press **Enter**.
2. Log in with the IBM i service tools user ID.
3. Select the **Start a service tool** option.
4. Select the **Display/Alter/Dump** option.
5. Select the **Display/Alter storage** option.
6. Select the **Licensed Internal Code (LIC) data** option.
7. Select the **Advanced analysis** option.
8. Type 1 in the **option** field at the top of the list.
9. Enter **VIOFR** in the corresponding **command** field and then press **Enter**.
10. Enter the **VIOFR** command options for the specific operations that are related to the virtual console and then press **Enter**. For more information, see [VIOFR command options for the virtual console](#).

To run the Advanced Analysis **VIOFR** command from DST, complete the following steps:

1. In the **5250 session DST** menu, select the **Start a service tool** option.
2. Select the **Display/Alter/Dump** option.
3. Select the **Display/Alter storage** option.
4. Select the **Licensed Internal Code (LIC) data** option.
5. Select the **Advanced analysis** option.
6. Type 1 in the **option** field at the top of the list.
7. Enter **VIOFR** in the corresponding **command** field and then press **Enter**.
8. Enter the **VIOFR** command options for the specific operations that are related to the virtual console and then press **Enter**. For more information, see [VIOFR command options for the virtual console](#).

VIOFR command options for the virtual console

1. To check the current security mode, type the following option:

```
Options: -tty
```

2. To enable the secure mode, type the following option:

```
Options: -tty -securetty
```

3. To disable the secure mode, type the following option:

```
Options: -tty -securetty disable
```

Note: When you switch modes, all active console sessions are closed to ensure a clean transition. Reconnect the console after the mode change completes.

Other security considerations

Configure a Transmission Control Protocol (TCP) port restriction to assign port 2301 exclusively to the System Service Representative (QSRV) users and prevent unauthorized applications from using it. To configure a TCP port restriction, complete either of the following steps:

Configuring a TCP port restriction by using the **CL** command:

```
ADDTCPPORT PORT(2301) PROTOCOL(*TCP) USRPRF(QSRV)
```

Configuring a TCP port restriction by using IBM i:

1. Navigate to **Network > Servers > TCP/IP**.
2. Select **Configuration > TCP/IP Configuration Properties**.
3. Enter the values in the following fields:

Table 18. Port restriction values	
Field	Supported values
Port	2301
Protocol	TCP
User Profile	QSRV

4. Click **Save**.

Starting the network-server description for an AIX logical partition

You can start the network-server description (NWSD) for an AIX logical partition that uses IBM i resources to make the resources defined in the NWSD available to the AIX logical partition.

About this task

To start (vary on) the NWSD for an AIX logical partition, complete the following tasks:

Procedure

1. If you use IBM Navigator for i, start the NWSD by using IBM Navigator for i.
 - a) Click **Network > Windows Administration > Integrated xSeries Servers**
 - b) Right-click the name of the NWSD that you want to start.
 - c) Click **Start**.
2. If you use the character-based interface, start the NWSD by using the character-based interface:
 - a) Type `WRKCFGSTS *NWS` and press Enter.
 - b) Type **1** next to the NWSD that you want to start and press Enter.

Creating an IBM i logical partition that uses IBM i virtual I/O resources

You can create an IBM i logical partition that uses IBM i virtual I/O resources on servers that are managed by a Hardware Management Console (HMC). This allows you to maximize utilization of the physical hardware and simplify the backup procedure for your managed system.

Before you begin

To set up an IBM i logical partition that uses IBM i virtual I/O resources, configure the following items:

- You must create a virtual server SCSI adapter for the IBM i logical partition that provides virtual SCSI disk resources, and create a virtual client SCSI adapter for the IBM i that uses virtual SCSI disk

resources. You can then set up the IBM i logical partition with the virtual server SCSI adapter to provide disk resources to the IBM i logical partition with the virtual client SCSI adapter through the virtual SCSI connection.

- You can create a virtual serial connection between the IBM i logical partition that provides the virtual resources and the IBM i logical partition that uses the virtual resources. A virtual serial connection allows you to connect to the IBM i logical partition that uses the virtual resources from the IBM i logical partition that provides the virtual resources.
- If you want to use virtual Ethernet, create two virtual Ethernet adapters on the IBM i logical partition that uses virtual I/O resources. Both virtual Ethernet adapters must be set to connect to a virtual Ethernet adapter on the IBM i logical partition that provides the virtual I/O resources. In other words, all three virtual Ethernet adapters must be set to the same virtual LAN ID.

Both logical partitions must use IBM i 6.1.1, or later.

Alternatively, you can create a Virtual I/O Server logical partition and configure the IBM i logical partition to use the virtual SCSI and virtual Ethernet resources of the Virtual I/O Server logical partition. You might need to enter a PowerVM Editions activation code to create a Virtual I/O Server logical partition on your server.

About this task

For more information about creating a logical partition, see “Creating logical partitions” on page 79. For more information about adding IBM i hosted virtual SCSI adapters when the HMC is at version 8.7.0, or later, see [Managing virtual storage for a partition](#).

Creating a network-server description and a network-server storage space for an IBM i logical partition that uses IBM i resources

A *network-server description (NWSD)* is an IBM i object that describes the storage resources that are used by an integrated operating environment. An NWSD can be linked to one or more network-server storage spaces. Create an NWSD to assign storage to an IBM i logical partition that uses IBM i virtual I/O resources.

Before you begin

If you assign multiple NWSDs to an IBM i logical partition that uses IBM i virtual I/O resources, ensure that only one of those NWSDs is set to provide virtual optical disk resources. Restrict the optical devices on all of the other NWSDs by adding RSTDDEVRS (*ALLOPT) to the CRTNWSD parameters for those NWSDs.

About this task

To create an NWSD and a network-server storage space for an IBM i logical partition that uses IBM i resources, follow these steps:

Procedure

1. Determine the correct SCSI server resource name.
 - If there is only one SCSI server adapter corresponding to a given client logical partition, and that adapter has its remote logical partition and remote slot configured correctly, you can specify *AUTO as the RSRCNAME in your NWSD.
 - Otherwise, you must determine the actual resource name. At an IBM i command line on the IBM i logical partition that provides virtual resources, type WRKHDWRSC *CMN, and find a controller resource with type 290B and a converged location code that corresponds to the SCSI server adapter at the Hardware Management Console (HMC). This resource name will be used later to specify the SCSI server resource.
2. At an IBM i command line on the logical partition that provides resources, type CRTNWSD to create a network server description, and press F4 for prompts.

3. Specify the following information:

The default or suggested parameter values are provided within the parentheses. These settings are relevant only to a logical partition. After the installation, if your root file system (/) is not installed on the first partition of the first disk, you must set a root parameter.

- NWSD (Provide a name for the NWSD)
- RSRNAME (*AUTO or the resource name of the SCSI server resource)
- TYPE(*GUEST *OPSYS)
- ONLINE (*NO or *YES)
- PARTITION ('Provide the name of your IBM i logical partition that uses IBM i resources')

As an alternative to the Partition parameter, you can also specify a logical partition number by typing PTNNBR(*integer*) where *integer* is the number of the logical partition you are specifying.

- CODEPAGE (437)
- TCPPTCFG (*NONE)
- RSTDDEVRSC (for virtual CD and tape devices) (*NONE for the NWSD that is to provide virtual optical disc resources, *ALLOPT for all other NWSDs)
- SYNCTIME (*TYPE)
- IPLSRC (*NWSSTG)
 - You can store a kernel in a disk partition of a virtual disk (a network-server storage space (NWSSTG)). By specifying the IPLSRC (*NWSSTG) parameter, you are specifying that the IBM i logical partition will start from a disk partition on that virtual disk. The disk partition on the virtual disk must be formatted as type PReP Boot (type 0x41) and marked as a device that starts.
 - To start an NWSD with a load source from a stream file, set the IPLSRC parameter to *STMF and set the IPLSTMF parameter to point to the load source. You must have read access to the file and the path leading to the file to use the vary on command.
- IPLSTMF (*NONE)
- IPLPARM (*NONE)
- PWRCTL (*YES)
 - If you specify PWRCTL (*YES), perform the following steps:
 - a. Ensure that the server adapter in the IBM i logical partition that provides the virtual resources specifies the remote logical partition and remote slot in its configuration.
 - b. Ensure that the client logical partition has the IBM i logical partition that provides the virtual resources as the power-controlling logical partition in the profile.
 - c. Ensure before you activate the NWSD that the partition profile for the IBM i logical partition that uses IBM i virtual I/O resources has been saved by activating the logical partition from the HMC, even if the logical partition does not activate correctly because of the absence of virtual devices.
 - If you specify PWRCTL (*NO), virtual devices will be available to the logical partition. You must shut down and restart the logical partition by using the HMC.

4. Create the network-server storage space using the interface that you prefer.

Interface	Actions
IBM Navigator for i	a. Expand My Connections > your server > Network > Windows Administration. b. Right-click the Disk Drives and select New Disk. c. In the Disk drive name field, specify the name that you want to give to the disk drive. d. In the Description field, specify a meaningful description for the disk drive.

Interface	Actions
	e. In the Capacity field, specify the size of the new disk drive in megabytes. f. Click OK.
IBM i character-based interface	a. At an IBM i command line on the IBM i logical partition that provides the virtual I/O resources, type the command CRTNWSSTG and press F4. The Create NWS Storage Space (CRTNWSSTG) display opens. b. In the Network-server storage space field, specify the name you want to give to the storage space. c. In the Size field, specify the size in megabytes for the new storage space. d. In the Format field, specify *OPEN. e. In the Text description field, specify a meaningful description for the storage space. f. Press Enter.

5. List the network server storage spaces on the logical partition by using the Work with Network Server Storage Spaces command.
6. If you use IBM Navigator for i, link the network-server storage space by using IBM Navigator for i.

Interface	Actions
IBM Navigator for i	a. Expand My Connections > your server > Network > Windows Administration. b. Click Disk Drives, right-click an available network-server storage space, and select Add Link. c. Select the server to which you want to link the network-server storage space. d. Select the link sequence position you want to use. e. Select one of the available data access types. f. Click OK.
IBM i character-based interface	a. At an IBM i command line on the IBM i logical partition that provides the virtual I/O resources, type the command ADDNWSSTGL and press F4. The Add Network-Server Storage Link (ADDNWSSTGL) display opens. b. In the Network server description field, specify the name of the network server description (NWSD). c. In the Dynamic storage link field, specify *YES to make the network-server storage space dynamically available to the IBM i logical partition that uses IBM i virtual I/O resources (that is, available without restarting the IBM i logical partition). d. In the Drive sequence number field, specify the link sequence position you want to use. e. Press Enter.

Connecting an IBM i logical partition that uses IBM i virtual I/O resources to the virtual console

You can connect an IBM i logical partition that uses IBM i virtual I/O resources to the virtual console so that you can install the operating system or access the command line interface for the IBM i logical partition that uses IBM i virtual I/O resources.

About this task

The virtual console provides the console function for an IBM i logical partition that uses IBM i virtual I/O resources. It is used primarily during the initial installation of the operating system. The virtual console can also be used to view server errors or to restore communication to the LAN. This console connection is used prior to configuring TCP/IP.

To open a terminal window, complete the following steps :



1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. Select the logical partition and click **Actions > Console > Open Dedicated 5250 Console**.

Starting the network-server description for an IBM i logical partition that uses IBM i virtual I/O resources

You can start the network-server description (NWSD) for an IBM i logical partition that uses IBM i resources to make the resources defined in the NWSD available to the IBM i logical partition that uses IBM i resources.

About this task

To start (vary on) the NWSD for an IBM i logical partition that uses IBM i resources, complete the following tasks:

Procedure

1. If you use IBM Navigator for i, start the NWSD by using IBM Navigator for i.
 - a) Click **Network > Windows Administration > Integrated xSeries Servers**
 - b) Right-click the name of the NWSD that you want to start.
 - c) Click **Start**.
2. If you use the character-based interface, start the NWSD by using the character-based interface:
 - a) Type `WRKCFGSTS *NWS` and press Enter.
 - b) Type **1** next to the NWSD that you want to start and press Enter.

Creating a Linux logical partition that uses IBM i virtual I/O resources

You can create a Linux logical partition that uses IBM i virtual I/O resources on servers that are managed by a Hardware Management Console (HMC). This allows you to maximize utilization of the physical hardware and simplify the backup procedure for your managed system.

Before you begin

To set this up, you must create virtual SCSI adapters that connect the logical partitions with each other. You can then set up the IBM i logical partition to provide disk resources to the Linux logical partition through the virtual SCSI connection. You can also create a virtual serial connection between the IBM i logical partition and the Linux logical partition. A virtual serial connection allows you to connect to the Linux logical partition from the IBM i logical partition.

Alternatively, you can create a Virtual I/O Server logical partition and configure the Linux logical partition to use the virtual SCSI and virtual Ethernet resources of the Virtual I/O Server logical partition. You might need to enter a PowerVM Editions activation code to create a Virtual I/O Server logical partition on your server.

For more information about creating a logical partition, see [“Creating logical partitions”](#) on page 79. For more information about adding IBM i hosted virtual SCSI adapters when the HMC is at version 8.7.0, or later, see [Managing virtual storage for a partition](#).

Creating an NWSD and a network-server storage space for a Linux logical partition

A *network-server description* (NWSD) is an IBM i object that describes the storage resources that are used by an integrated operating environment. An NWSD can be linked to one or more network-server storage spaces. Create an NWSD to assign storage to a Linux logical partition that uses IBM i resources.

About this task

To create an NWSD and a network-server storage space for a Linux logical partition that uses IBM i resources, follow these steps:

Procedure

1. Determine the correct SCSI server resource name.

- If there is only one SCSI server adapter corresponding to a given client logical partition, and that adapter has its remote logical partition and remote slot configured correctly, you can specify *AUTO as the RSRNAME in your NWSD.
- Otherwise, you must determine the actual resource name. At an IBM i command line, type WRKHDWRSC *CMN, and find a controller resource with type 290B and a converged location code that corresponds to the SCSI server adapter at the Hardware Management Console (HMC). This resource name will be used later to specify the SCSI server resource.

2. At an IBM i command line on the logical partition that shares resources, type CRTNWSD and press F4 for prompts.

3. Specify the following information.

The default or suggested parameter values are provided within the parentheses. These settings are relevant only to a logical partition. After the installation, if your root file system (/) is not installed on the first partition of the first disk, you must set a root parameter.

- NWSD (Provide a name for the NWSD)
- RSRNAME (*AUTO or the resource name of the SCSI server resource)
- TYPE(*GUEST)
- ONLINE (*NO or *YES)
- PARTITION ('Provide the name of your Linux logical partition')

As an alternative to the Partition parameter, you can also specify a logical partition number by typing PTNNBR(*integer*) where *integer* is the number of the logical partition you are specifying.

- CODEPAGE (437)
- TCPPOPTCFG (*NONE)
- RSTDDEVRSC (for virtual CD and tape devices) (*NONE)
- SYNCTIME (*TYPE)
- IPLSRC (*NWSSTG)
 - You can store a kernel in a disk partition of a virtual disk (a network-server storage space (NWSSTG)). By specifying the IPLSRC (*NWSSTG) parameter, you are specifying that the Linux logical partition will start from a disk partition on that virtual disk. The disk partition on the virtual disk must be formatted as type PReP Boot (type 0x41) and marked as a device that starts. You can format a disk partition as type PReP Boot by using the Linux **fdisk** command with the -t

- option. You can specify that the disk partition starts by using the **fdisk** command with the -a option.
- To start an NWSD with a kernel from a stream file, set the IPLSRC parameter to *STMF and set the IPLSTMF parameter to point to the kernel. You must have read access to the file and the path leading to the file to use the vary on command. This value only loads the kernel. After the kernel is running, it must find a root file system. In an initial installation, the root file system might be a RAM disk that is physically attached to the kernel.
 - IPLSTMF (*NONE)
 - IPLPARM (*NONE)
 - PWRCTL (*YES)
 - If you specify PWRCTL (*YES), perform the following steps:
 - a. Ensure that the server adapter in the IBM i logical partition specifies the remote logical partition and remote slot in its configuration.
 - b. Ensure that the client logical partition has the IBM i logical partition as the power-controlling logical partition in the profile.
 - c. Ensure before you activate the NWSD that the client logical partition's profile has been saved to the server by activating the logical partition from the HMC, even if the client operating system does not activate correctly because of the absence of virtual devices.
 - If you specify PWRCTL (*NO), virtual devices will be available to the logical partition. You must shut down and restart the logical partition using the HMC.
4. If you use IBM Navigator for i, create the network-server storage space by using IBM Navigator for i.
 - a) Expand **My Connections > your server > Network > Windows Administration**.
 - b) Right-click the **Disk Drives** and select **New Disk**.
 - c) In the **Disk drive name** field, specify the name that you want to give to the disk drive.
 - d) In the **Description** field, specify a meaningful description for the disk drive.
 - e) In the **Capacity** field, specify the size of the new disk drive in megabytes.
Refer to your preferred Linux distributor installation documentation to determine the size you want to use.
 - f) Click **OK**.
 - g) Continue with step “6” on page 106.
 5. If you use a character-based interface, create the network-server storage space using the character-based interface:
 - a) At an IBM i command line, type the command CRTNWSSTG and press F4.
The Create NWS Storage Space (CRTNWSSTG) display opens.
 - b) In the Network-server storage space field, specify the name you want to give to the storage space.
 - c) In the Size field, specify the size in megabytes for the new storage space.
Refer to your preferred Linux distributor installation documentation to determine the size you want to use.
 - d) In the Text description field, specify a meaningful description for the storage space.
 - e) Press Enter.
 - f) Continue with step “7” on page 107.
 6. If you use IBM Navigator for i, link the network-server storage space by using IBM Navigator for i:
 - a) Expand **My Connections > your server > Network > Windows Administration**.
 - b) Click **Disk Drives**, right-click an available network-server storage space, and select **Add Link**.
 - c) Select the server to which you want to link the network-server storage space.
 - d) Select the link sequence position you want to use.
 - e) Select one of the available data access types.

f) Click **OK**.

The procedure is complete. Do not complete step “7” on page 107.

7. If you use a character-based interface, link the network-server storage space by using the character-based interface:
 - a) At an IBM i command line, type the command ADDNWSSTGL and press F4.
The Add Network-Server Storage Link (ADDNWSSTGL) display opens.
 - b) In the Network server description field, specify the name of the network server description (NWSD).
 - c) In the Dynamic storage link field, specify *YES to make the network-server storage space dynamically available to the logical partition (that is, available without rebooting the Linux logical partition).
 - d) In the Drive sequence number field, specify the link sequence position you want to use. If you want the system to find the next available position for you, specify *CALC.
 - e) Press Enter.

Connecting to the virtual console for a Linux logical partition

You can connect to the virtual console for a Linux logical partition so that you can install the operating system or access the command line interface for the Linux logical partition.

Before you begin

You must have one of the following privileges to use the Linux virtual console.

- System Partitions - Operations
- System Partitions - Administration

About this task

The virtual console provides console functions for a Linux server. Use it primarily during initial operating system installation. You can also use the virtual console to view server errors or restore local area network (LAN) communication. Use this console connection before you configure the Transmission Control Protocol (TCP) or Internet Protocol (IP).

You can use any Telnet client as the Linux console. Multiple Telnet clients can share access to the same virtual console. TCP/IP must be configured and running on at least one IBM i logical partition.

For IBM i version 7.4, IBM i version 7.5 Technology Refresh 8, and IBM i version 7.6 Technology Refresh 2, the virtual console server supports the following operational modes:

- Unsecured mode: The server accepts direct Telnet connections on port 2301 from any network interface. The unsecured mode is the default setting for IBM i version 7.4, IBM i version 7.5, and IBM i version 7.6.
- Secure mode: The server accepts connections only from the localhost 127.0.0.1:2301 port. To access the server remotely, create an SSH tunnel with the local port forwarding capability, and then connect with a Telnet client to the forwarded local port. SSH encrypts all traffic, including login credentials.

Default security modes by IBM i release:

- Version 7.4: Unsecured mode
- Version 7.5 Technology Refresh 8: Unsecured mode
- Version 7.6 Technology Refresh 2: Unsecured mode

Using secure mode to connect to the virtual console

In the secure mode, create an SSH tunnel before you connect to the virtual console. The SSH tunnel encrypts all traffic, including service tools credentials.

Prerequisites and privileges for secure mode:

- OpenSSH is installed and the OpenSSH server is running on IBM i. Install OpenSSH with license program 5733-SC1 Option 1.
- PASE must be installed. Install PASE with 5770-SS1 Option 33.
- The SSH client is available on the workstation and installed in AIX, Linux, Mac, and Windows version 10 or 11
- IBM i users must have the Secure Shell (SSH) access
- Service tool user ID must have System partitions - operations or System partitions - administration access.
- For the IBM i version 7.4, the PTF MJ09039 must be installed for security mode changes
- For IBM i version 7.5 Technology Refresh 8, the PTF MJ08967 must be installed for security mode changes
- For IBM i version 7.6 Technology Refresh 2, the PTF MJ08968 must be installed for security mode changes

To connect to a virtual console through the secure mode, complete the following steps:

1. Establish the SSH tunnel:

- a. Open the console or the command line on your workstation and use the local port forwarding process to create an SSH tunnel. An example command using a profile with the SSH access follows:

```
ssh -L localhost:23011:localhost:2301 IBMiUserWithSSHAccess@IBMiHostname
```

- b. Enter the password of the IBM i profile in step 1a.

Note: Do not close the SSH session.

- c. Follow step 2 to connect with a standard Telnet Client, or step 3 to connect with the IBM Personal Communications client.

2. Connect with a standard Telnet client:

- a. Open a second console or command line on your workstation, separate from the SSH tunnel, and connect your Telnet client to the local forwarded port you selected in Step 1a. An example command follows:

```
telnet localhost <your_chosen_port>
```

Note: In the example in step 1a, the **your_chosen_port** value is 23011.

- b. Select the logical partition for which you want to connect to the virtual console.
- c. Enter the IBM i service tools ID and password.

Note: If Multi-Factor Authentication (MFA) is configured for the service tools ID, enter the One-Time Password (TOTP) when prompted.

- d. The telnet client is now securely connected to the console of the selected logical partition. After you finish using the console, close the Telnet session and then enter `exit` in the SSH console to close the SSH connection. This completes the procedure of connecting to a virtual console through the secure mode.

3. Connect with the IBM Personal Communications:

- a. Click **Start > IBM Personal Communications > Start or Configure Session**. The **Customize Communication** page is displayed.
- b. Select **ASCII** as your type of host and select **Link Parameters**. The **Telnet ASCII** page is displayed.
- c. Enter the hostname or the IP address of the logical partition that is sharing its resources, and enter the local port number that you chose in step 1a.

Note: In the example in step 1a, the **your_chosen_port** value is 23011.

- d. Click **OK**.

- e. From the **Guest Partition Consoles** page, select the logical partition to which you want to connect as the console.
- f. Enter the IBM i service tools ID and password to connect to the Linux logical partition.

Note: If Multi-Factor Authentication (MFA) is configured for the service tools ID, enter the One-Time Password (TOTP) when prompted.

- g. The IBM Personal Communications telnet client is now securely connected to the console of the selected logical partition. After you finish using the console, close the Telnet session and then enter `exit` in the SSH console to close the SSH connection. This completes the procedure of connecting to a virtual console through the secure mode.

Note: When you run the local x5250 virtual console session with the **TELNET** Command Language (CL) command to connect to the local host system, the console displays all typed characters on the screen, including service tools credentials. This behavior creates a security risk during screen share sessions where other users can view your credentials. To mitigate this risk, use a separate SSH tunnel with a standard telnet client or IBM Personal Communications telnet client, even when you connect from the local host system.

Using unsecured mode to connect to the virtual console

The unsecured mode transmits IBM i credentials and console traffic in an unencrypted form across the network. Use this mode only on local networks where security is not a concern.

Note: In IBM i versions before 7.4 with PTF MJ09039, 7.5 Technology Refresh 8, and 7.6 Technology Refresh 2 versions, the console server accepts telnet connections from all network interfaces, including the `localhost`. You can still use the SSH tunnel method to establish a secure connection by encrypting the telnet login credentials and the connection data on the network. For more information see [Using secure mode to connect to the virtual console](#).

To connect to a virtual console through the unsecured mode, complete the following steps:

Note: Follow step 1 to connect with a standard Telnet Client, or step 2 to connect with IBM Personal Communications.

1. Connect with a standard Telnet client:

- a. Open a console or command prompt on your workstation.
- b. Run the **Telnet** command to connect your IBM i system using port 2301. Example:

```
telnet IBMiHostname 2301
```

Note: The `IBMiHostname` is the IBM i system hostname or IP address.

- c. Select the logical partition for which you want to connect to the virtual console.
- d. Enter the IBM i service tools ID and password.

Note: If Multi-Factor Authentication (MFA) is configured for the service tools ID, enter the One-Time Password (TOTP) when prompted.

- e. The telnet client is now connected to the console of the selected logical partition. After you finish using the console, close the Telnet session. This completes the procedure of connecting to a virtual console through the unsecure mode.

2. Connect with the IBM Personal Communications:

- a. Click **Start > IBM Personal Communications > Start or Configure Session**. The **Customize Communication** page is displayed.
- b. Select **ASCII** as your type of host and select **Link Parameters**. The **Telnet ASCII** page is displayed.
- c. Enter the hostname or the IP address of the logical partition that is sharing its resources, and enter port number 2301.
- d. Click **OK**.
- e. From the **Guest Partition Consoles** page, select the logical partition to which you want to connect as the console.

f. Enter the IBM i service tools ID and password.

Note: If Multi-Factor Authentication (MFA) is configured for the service tools ID, enter the One-Time Password (TOTP) when prompted.

g. The IBM Personal Communications telnet client is now connected to the console of the selected logical partition. After you finish using the console, close the Telnet session. This completes the procedure of connecting to a virtual console through the unsecure mode.

Managing the virtual console security mode

Use the **VIOFR** command to manage virtual console security mode and switch between secure and unsecured modes without an Initial Program Load (IPL). Run the **VIOFR** command from System Service Tools (SST) or Dedicated Service Tools (DST).

You must have the following privileges to run the **VIOFR** command:

- Service tools user ID
- Display/Alter/Dump privilege
- Start Service tool privilege
- Service tool security privilege

To run the **VIOFR** command from SST, complete the following steps:

1. In the 5250 command session, enter STRSST and press **Enter**.
2. Log in with the IBM i service tools user ID.
3. Select the **Start a service tool** option.
4. Select the **Display/Alter/Dump** option.
5. Select the **Display/Alter storage** option.
6. Select the **Licensed Internal Code (LIC) data** option.
7. Select the **Advanced analysis** option.
8. Type 1 in the **option** field at the top of the list.
9. Enter **VIOFR** in the corresponding **command** field and then press **Enter**.
10. Enter the **VIOFR** command options for the specific operations that are related to the virtual console and then press **Enter**. For more information, see [VIOFR command options for the virtual console](#).

To run the Advanced Analysis **VIOFR** command from DST, complete the following steps:

1. In the **5250 session DST** menu, select the **Start a service tool** option.
2. Select the **Display/Alter/Dump** option.
3. Select the **Display/Alter storage** option.
4. Select the **Licensed Internal Code (LIC) data** option.
5. Select the **Advanced analysis** option.
6. Type 1 in the **option** field at the top of the list.
7. Enter **VIOFR** in the corresponding **command** field and then press **Enter**.
8. Enter the **VIOFR** command options for the specific operations that are related to the virtual console and then press **Enter**. For more information, see [VIOFR command options for the virtual console](#).

VIOFR command options for the virtual console

1. To check the current security mode, type the following option:

```
Options: -tty
```

2. To enable the secure mode, type the following option:

```
Options: -tty -securetty
```

3. To disable the secure mode, type the following option:

```
Options: -tty -securetty disable
```

Note: When you switch modes, all active console sessions are closed to ensure a clean transition. Reconnect the console after the mode change completes.

Other security considerations

Configure a Transmission Control Protocol (TCP) port restriction to assign port 2301 exclusively to the System Service Representative (QSRV) users and prevent unauthorized applications from using it. To configure a TCP port restriction, complete either of the following steps:

Configuring a TCP port restriction by using the **CL** command:

```
ADDTCPPORT PORT(2301) PROTOCOL(*TCP) USRPRF(QSRV)
```

Configuring a TCP port restriction by using IBM i:

1. Navigate to **Network > Servers > TCP/IP**.
2. Select **Configuration > TCP/IP Configuration Properties**.
3. Enter the values in the following fields:

Table 19. Port restriction values	
Field	Supported values
Port	2301
Protocol	TCP
User Profile	QSRV

4. Click **Save**.

Starting the network-server description for a Linux logical partition

You can start the network-server description (NWSD) for a Linux logical partition that uses IBM i resources to make the resources defined in the NWSD available to the Linux logical partition.

About this task

To start (vary on) the NWSD for a Linux logical partition, complete the following steps:

Procedure

1. If you use IBM Navigator for i, start the NWSD by using IBM Navigator for i.
 - a) Click **Network > Windows Administration > Integrated xSeries Servers**
 - b) Right-click the name of the NWSD that you want to start.
 - c) Click **Start**.
2. If you use the character-based interface, start the NWSD by using a character-based interface:
 - a) Type `WRKCFGSTS *NWS` and press Enter.
 - b) Type **1** next to the NWSD that you want to start and press Enter.

Designating the service logical partition for your managed system

The *service logical partition* is the IBM i logical partition on a server that you can configure to apply server firmware updates to the service processor or to the hypervisor. You can also use the service logical partition to communicate server common hardware errors to IBM. These abilities are useful if the Hardware Management Console (HMC) is undergoing maintenance or is otherwise unable to perform these functions.

Before you begin

The preferred method for applying server firmware updates and communicating server common hardware errors to IBM is by using the HMC.

Servers that do not have IBM i logical partitions also do not have a service logical partition. If these servers are managed by an HMC, then you must use the HMC to update the server firmware, and the servers can contact service and support only through the HMC. Use a backup HMC to ensure that the servers have redundant methods for contacting service and support and for applying fixes.

You can designate only one logical partition at a time as the service logical partition for your managed system. The service logical partition for your server must be an IBM i logical partition.

Before you can designate a logical partition as the service logical partition for your managed system, you must shut down the logical partition. You must also shut down the logical partition before you remove the service logical partition designation from the logical partition. If you want to change the service logical partition from one logical partition to another logical partition, you must shut down both logical partitions before using this procedure.

Note: You must designate a service logical partition on a server only after you use the HMC to create, change, delete, copy, or activate any logical partitions on the managed system. You can set up the operating system on an unpartitioned server to contact service and support, and you can use the operating system on an unpartitioned server to apply server firmware updates.

About this task

To designate one of your logical partitions as the service logical partition, complete the following steps:



1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **Actions** > **View System Properties**. The **Properties** page is displayed.
4. In the **General Settings** area, select the logical partition that you want to designate as the service logical partition.
5. Click **OK**.

Resetting the managed system to a nonpartitioned configuration

You can use the Hardware Management Console (HMC) and the Advanced System Management Interface (ASMI) to erase all of your logical partitions and reset the managed system to a nonpartitioned configuration. When you reset the managed system, all of the physical hardware resources are assigned to a single logical partition. This allows you to use the managed system as if it were a single, nonpartitioned server.

Before you begin



Attention: By resetting a partitioned managed system to a nonpartitioned configuration, you will lose all of your logical partition configuration data. However, resetting the managed system does not erase the operating systems and data from disk units on that managed system.

Before you reset the managed system, ensure that the hardware placement in the managed system supports a nonpartitioned configuration. If the hardware placement in the managed system does not support a nonpartitioned configuration, you must move the hardware so that the hardware placement supports a nonpartitioned configuration. For more information about how to place the hardware in your managed system to support a nonpartitioned configuration, contact your marketing representative or business partner.

Also, if you plan to use an operating system that is already installed on one of the logical partitions on the managed system (instead of reinstalling the operating system after you reset the managed system), consider how the console used by that operating system will change when you reset the managed system. If the operating system that you want to use is AIX, log into AIX and enable the login prompt for the virtual serial port vty0 using either the System Management Interface Tool (SMIT) or the **chdev** command. You can then reset the managed system, use a physical serial console to log into AIX, and use SMIT or the **chcons** command to change the console device to the console device you want to use.

You must have an ASMI login profile with an administrator authority level.

Parts of this procedure must be performed *at your HMC* (not connected remotely). Ensure that you have physical access to the HMC before you begin.

About this task

To reset a managed system with logical partitions to a nonpartitioned configuration using the HMC, follow these steps:

Procedure

1. Shut down all logical partitions on your managed system using operating system procedures.
For more information about shutting down logical partitions using operating system procedures, see the following information:
 - For logical partitions running AIX, see [“Shutting down AIX logical partitions” on page 137](#).
 - For logical partitions running Linux, see [“Shutting down Linux logical partitions” on page 143](#).
 - For logical partitions running Virtual I/O Server, see [“Shutting down Virtual I/O Server logical partitions by using the HMC” on page 144](#).
2. If the managed system powered off automatically when you shut down the last logical partition, power on the managed system to the Standby state. Complete the following steps:
 - a) In the navigation pane of your HMC, open **Systems Management** and click **Servers**.
 - b) In the work pane, select the managed system, click the **Tasks** button, and click **Operations > Power On**.
 - c) Select the power-on mode of **Partition Standby** and click **OK**.
 - d) Wait until the work pane displays a Standby state for the managed system.
3. Initialize the profile data on the HMC. Complete the following:
 - a) In the work pane, select the managed system, click the **Tasks** button, and click **Configuration > Manage Partition Data > Initialize**.
 - b) Click **Yes** to confirm.
4. Clear the logical partition configuration data on the managed system. Complete the following *at your HMC* (not connected remotely):
 - a) In the navigation pane, click **HMC Management**.
 - b) In the work pane, click **Open Restricted Shell Terminal**.
 - c) Type the command: `lpcfgop -m managed_system_name -o clear`, where *managed_system_name* is the name of the managed system as it is displayed in the work pane.
 - d) Enter **1** to confirm.
This step takes several seconds to complete.
5. Optional: If you no longer intend to manage the system using the HMC, remove the connection between the HMC and the managed system. To remove the connection between the HMC and the managed system, complete the following:
 - a) In the work pane, select the managed system, click the **Tasks** button, and click **Connections > Reset or Remove Connection**.
 - b) Select **Remove connection** and click **OK**.

6. Access the Advanced System Management Interface (ASMI) using a Web browser on a PC. If you do not already have a PC that is set up to access the ASMI on the managed system, you need to set up the PC at this point.

For instructions, see [Accessing the ASMI using a web browser](#).

7. On the ASMI Welcome pane, log in using the admin user ID (enter admin into **User ID**, enter the admin password into **Password**, and click **Log In**).
8. In the navigation pane, expand **Power/Restart Control** and click **Power On/Off System**.
9. Set **Boot to server firmware** to Running.
10. Click **Save settings and power off**.
11. Click **Power On/Off System** periodically to refresh the window. Repeat this step until **Current system power state: Off** is displayed in the navigation pane.
12. Click **Save settings and power on**.
13. Wait for the managed system to restart.

It can take several minutes for the managed system and operating system to restart completely.

Deleting a logical partition

You can use the Hardware Management Console (HMC) to delete a logical partition and all of the partition profiles associated with the logical partition.

Before you begin

You cannot delete a logical partition if it is the service logical partition of your managed system. Before you can delete such a logical partition, you must designate another logical partition as the service logical partition of your managed system or remove the service logical partition designation from the logical partition.

Before you delete a logical partition, complete the following steps:

1. Shut down the logical partition that you plan to delete. For instructions, see [“Shutting down and restarting logical partitions” on page 137](#).
2. If the logical partition that you plan to delete is a Virtual I/O Server logical partition that is assigned to the shared memory pool (hereafter referred to as a *paging VIOS partition*), remove the paging VIOS partition from the shared memory pool. For instructions, see [“Removing a paging VIOS partition from the shared memory pool” on page 129](#).

When you delete a logical partition that uses shared memory, the HMC automatically performs the following tasks:

- The HMC removes the shared memory partition from the shared memory pool.
- The HMC returns the physical memory that was allocated to the shared memory partition for its I/O devices to the shared memory pool so that the hypervisor can allocate the physical memory to other shared memory partitions.
- The HMC releases the paging space device that was allocated to the shared memory partition so that it becomes available for other shared memory partitions to use.



Attention: This procedure erases the logical partition and the logical partition configuration data stored on the partition profiles.

About this task

To delete a logical partition, complete the following steps:



1. In the navigation pane, click the **Resources** icon.
2. Click **All Partitions**. The All Partitions page is displayed.

3. Select the logical partition and click **Tasks > Delete Partition**. You can select the **Cleanup associated Virtual I/O Server mappings** check box and the **Delete associated virtual disks** check box.
4. Click **OK** to confirm.

Configuring virtual resources for logical partitions

You can use the Hardware Management Console (HMC) to configure virtual resources such as virtual Ethernet adapters, Host Ethernet Adapter, and shared processor pools. Configure virtual resources to help optimize the use of physical system resources.

Note: HEA is not supported on POWER9 processor-based server.

Configuring Active Memory Expansion for AIX logical partitions

You can configure Active Memory Expansion for an AIX logical partition by using the Hardware Management Console (HMC). Configuring Active Memory Expansion for a logical partition compresses the memory of the logical partition and thus expands its memory capacity.

Before you begin

You can configure Active Memory Expansion for logical partitions that use dedicated memory and logical partitions that use shared memory.

Before you start, complete the following tasks:

1. Complete the required preparation tasks for Active Memory Expansion and ensure that your configuration meets the configuration requirements for Active Memory Expansion. For instructions, see [“Preparing to configure Active Memory Expansion” on page 64](#).
2. Enter the required activation code to enable Active Memory Expansion on the server. For instructions, see [“Entering the activation code for Active Memory Expansion” on page 124](#).

About this task

For more information about configuring Active Memory Expansion on a logical partition, see [Changing memory settings](#).

What to do next

After you configure Active Memory Expansion for the logical partition, monitor the performance of the workload and adjust the configuration, if necessary. For instructions, see [“Adjusting the Active Memory Expansion configuration to improve performance” on page 210](#).

Configuring a virtual Ethernet adapter

You can configure a virtual Ethernet adapter dynamically for a running logical partition by using the Hardware Management Console (HMC). Doing so will connect the logical partition to a virtual LAN (VLAN).

Before you begin

You can dynamically configure a virtual Ethernet adapter for a Linux logical partition only if the following conditions are met:

- A Linux distribution that supports dynamic partitioning is installed on the Linux logical partition. Distributions that support dynamic partitioning include SUSE Linux Enterprise Server 9 and later versions.
- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the [Service and productivity tools for Linux on POWER systems](#) website.

If you plan to configure an Ethernet adapter for a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*), you might need to adjust the amount of I/O entitled memory assigned to the shared memory partition before you configure the adapter:

- If the I/O entitled memory mode of the shared memory partition is set to the auto mode, you do not need to take action. When you configure the new Ethernet adapter, the HMC automatically increases the I/O entitled memory of the shared memory partition to accommodate the new adapter.
- If the I/O entitled memory mode of the shared memory partition is set to the manual mode, you must increase the I/O entitled memory that is assigned to the shared memory partition to accommodate the new adapter. For instructions, see [“Adding and removing I/O entitled memory dynamically to and from a shared memory partition”](#) on page 158.

About this task

For more information about managing virtual network connections on a logical partition when the HMC is at version 8.7.0, or later, see [Managing virtual network connections](#).

Note: In the AIX operating system, the receive buffers pools of the Virtual Ethernet increases in size and shrinks. When the load increases, the receive buffers pools increases by several buffers. The buffer pools can increase in size until the maximum value is reached (defined by the *buf_mode* attribute). When the load decreases, the receive buffers pools shrink to the minimum value (defined by the *buf_mode* attribute). The *buf_mode* attribute has the following possible values:

- *min* - Allocate the minimum buffer values. Increase as required and shrink back to min values.
- *max* - Allocate the maximum buffer values. Shrink is disabled. Fail the device open operation if the maximum values cannot be allocated.
- *max_min* - Attempt the maximum mode. If the maximum values cannot be allocated, then fall back to minimum mode.

You can run the **entstat** command on the Virtual Ethernet adapter to display the number of buffers that are allocated by the device.

What to do next

After you have finished, access any existing partition profiles for the logical partition and add the virtual Ethernet adapters to those partition profiles. The virtual Ethernet adapter is lost if you shut down the logical partition and activate that logical partition using a partition profile that does not have the virtual Ethernet adapter in it.

Changing the VLAN IDs of a virtual Ethernet adapter

You can dynamically change the VLAN IDs of a virtual Ethernet adapter for a running logical partition by using the Hardware Management Console (HMC).

Before you begin

Before you start, verify that the Virtual I/O Server is at version 2.2.0.0, or later.

About this task

For more information about managing virtual network connections on a logical partition when the HMC is at version 8.7.0, or later, see [Managing virtual network connections](#).

What to do next

After changing the VLAN IDs of a virtual Ethernet adapter, access any existing partition profiles for the logical partition and add the VLAN IDs of the virtual Ethernet adapters to the partition profiles. The dynamically set values in the additional VLANs are lost if you have shut down and then activated the logical partition by using a partition profile that does not have the new list of VLAN IDs of the virtual Ethernet adapter.

Configuring the Quality of Service priority for a virtual Ethernet adapter

You can dynamically configure the Quality of Service (QoS) priority of a virtual Ethernet adapter of a running logical partition by using the Hardware Management Console (HMC). You can prioritize the logical partition network traffic by specifying the value of IEEE 802.1Q priority level for each virtual Ethernet adapter.

About this task

For more information about managing virtual network connections on a logical partition when the HMC is at version 8.7.0, or later, see [Managing virtual network connections](#).

The virtual Ethernet QoS priority level values are in the range of 1 - 7. The following table lists the different priority levels.

VLAN user priority level	Quality of Service priority
1	Background
2	Spare
0 (default)	Best effort
3	Excellent effort
4	Controlled load
5	Video < 100 ms latency and jitter
6	Voice < 10 ms latency and jitter
7	Network control

MAC address controls using the HMC

The HMC version 7 release 7.2.0 or later, introduces HMC controls and policies for MAC address assignment to virtual Ethernet adapters and to logical Host Ethernet Adapter (LHEA).

By using the HMC, you can perform the following tasks:

- Specify a custom MAC address for the virtual Ethernet adapters of a logical partition.

Note: For a virtual Ethernet adapter, the default value is the HMC-generated MAC address.

Tip: Avoid specifying a MAC address to enable automatic generation of a MAC address.

- Apply the following controls to MAC address overrides specified at the operating system level:
 - Allow all operating-system-defined MAC addresses
 - Deny all operating-system-defined MAC addresses
 - Specify allowable operating-system-defined MAC addresses (you can specify a maximum of four operating-system defined MAC addresses)

Note: By default, all overrides are allowed. This is applicable to both virtual Ethernet adapters and LHEA. HEA is not supported on POWER9 processor-based server.

- Specify an optional initial MAC address for a virtual Ethernet adapter to replace an HMC-generated initial MAC address.

Note: The MAC address controls can be applied only when creating a logical partition, modifying a partition profile, or dynamically adding a virtual Ethernet adapter and logical Host Ethernet Adapter. You cannot dynamically modify an existing virtual Ethernet adapter or LHEA to add or change MAC controls.

The rules for custom virtual Ethernet MAC addresses are:

- The MAC address must be 6 bytes long.
- The bit 1 of byte 0 is reserved for Ethernet multicasting and must always be off.

- The bit 2 of byte 0 indicates that the MAC address is a locally administered address and must always be on.

Configuring the MAC address controls for a virtual Ethernet adapter

By using the Hardware Management Console (HMC), you can configure the MAC address controls of a virtual Ethernet adapter of a logical partition during logical partition creation, during partition profile modification, or when dynamically adding a virtual Ethernet adapter. You can also specify controls to MAC address overrides that are specified at the operating-system level.

About this task

For more information about managing virtual network connections on a logical partition when the HMC is at version 8.7.0, or later, see [Managing virtual network connections](#).

Configuring a virtual Fibre Channel adapter

You can configure a virtual Fibre Channel adapter dynamically for a running logical partition by using the Hardware Management Console (HMC).

Before you begin

A Linux logical partition supports the dynamic addition of virtual Fibre Channel adapters only if the DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the [Service and productivity tools for Linux on POWER systems](#) website.

If you plan to configure a virtual Fibre Channel adapter for a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*), you might need to adjust the amount of I/O entitled memory assigned to the shared memory partition before you configure the adapter:

- If the I/O entitled memory mode of the shared memory partition is set to the auto mode, you do not need to take action. When you configure the new virtual Fibre Channel adapter, the HMC automatically increases the I/O entitled memory of the shared memory partition to accommodate the new adapter.
- If the I/O entitled memory mode of the shared memory partition is set to the manual mode, you must increase the I/O entitled memory that is assigned to the shared memory partition to accommodate the new adapter. For instructions, see [“Adding and removing I/O entitled memory dynamically to and from a shared memory partition” on page 158](#).

When you dynamically configure a virtual Fibre Channel adapter on a client logical partition that uses Virtual I/O Server resources, the virtual Fibre Channel adapter is lost when you restart the logical partition because the partition profile does not include the virtual Fibre Channel adapter. You cannot add the virtual Fibre Channel adapter to a partition profile after you dynamically configure it on the logical partition because the virtual Fibre Channel adapter that you add to the partition profile is assigned a different pair of worldwide port names (WWPNs) than the virtual Fibre Channel adapter that you dynamically configured on the logical partition. If you want to include the virtual Fibre Channel adapter in a partition profile, then do not dynamically configure the virtual Fibre Channel adapter on the logical partition. Instead, create the virtual Fibre Channel adapter in a partition profile and then start the logical partition using that partition profile. For instructions, see [“Changing partition profile properties” on page 149](#).

About this task

For more information about assigning virtual Fibre Channel storage to a logical partition, see [Assigning virtual Fibre Channel storage to a partition](#).

What to do next

If you created a virtual Fibre Channel adapter on a Virtual I/O Server logical partition, complete the following tasks:

1. Access any existing partition profiles for the Virtual I/O Server logical partition and add the virtual Fibre Channel adapter to those partition profiles. The virtual Fibre Channel adapter is lost when you shut down the Virtual I/O Server logical partition and activate it by using a partition profile that does not include the virtual Fibre Channel adapter.
2. Assign the virtual Fibre Channel adapter to a physical port on the physical Fibre Channel adapter that is connected to the physical storage that you want the associated client logical partition to access. For instructions, see [Assigning the virtual Fibre Channel adapter to a physical Fibre Channel adapter](#).

Configuring physical ports on a Host Ethernet Adapter

You can use a Hardware Management Console (HMC) to configure the properties of each physical port on a Host Ethernet Adapter (HEA). These properties include port speed, duplex mode, maximum packet size, flow control setting, and the promiscuous logical partition for unicast packets. The physical port properties are also used by the logical ports that are associated with each physical port. HEAs are also known as Integrated Virtual Ethernet adapters (IVE adapters).

About this task

For more information about managing Host Ethernet Adapters on a logical partition when the HMC is at version 8.7.0, or later, see [Managing Host Ethernet Adapters](#).

What to do next

After this procedure is complete, you might need to reconfigure any logical ports that are associated with the changed physical ports. For example, if you change the maximum packet size on the physical port, you might also need to access the operating systems that use the resources on that physical port and change the maximum packet size for the corresponding logical ports.

Configuring shared processor pools

If your managed system supports more than one shared processor pool, you can use the Hardware Management Console (HMC) to configure shared processor pools on your managed system in addition to the default shared processor pool. These additional shared processor pools allow you to limit the processor usage of the logical partitions that belong to the shared processor pools. All shared processor pools other than the default shared processor pool must be configured before you can assign logical partitions to these shared processor pools.

Before you begin

You can use this procedure only if the managed system supports more than one shared processor pool and the HMC is at version 7 release 3.2.0, or later.

The default shared processor pool is preconfigured, and you cannot change the properties of the default shared processor pool.

About this task

For more information about managing shared processor pools, see [Managing shared processor pools](#).

What to do next

After this procedure is complete, assign logical partitions to the configured shared processor pools. You can assign a logical partition to a shared processor pool at the time that you create the logical partition, or you can reassign existing logical partitions from their current shared processor pools to the shared processor pools that you configured using this procedure. For instructions, see [“Reassigning logical partitions to shared processor pools” on page 125](#).

When you no longer want to use a shared processor pool, you can unconfigure the shared processor pool by using this procedure to set the maximum number of processing units and reserved number of

processing units to 0. Before you can unconfigure a shared processor pool, you must reassign all logical partitions that use the shared processor pool to other shared processor pools.

Configuring the shared memory pool

You can configure the size of the shared memory pool, assign paging space devices to the shared memory pool, and assign one or two Virtual I/O Server (VIOS) logical partitions (that provide access to the paging space devices) to the shared memory pool using the Hardware Management Console (HMC).

Before you begin

Before you start, complete the following tasks:

1. Enter the activation code for the PowerVM Enterprise Edition. For instructions, see [Entering the activation code for PowerVM Editions using the HMC version 7](#). The ability to share memory among multiple logical partitions is known as the PowerVM Active Memory Sharing technology. The PowerVM Active Memory Sharing technology is available with the PowerVM Enterprise Edition for which you must obtain and enter a PowerVM Editions activation code.
2. Ensure that your configuration meets the configuration requirements for shared memory. To review the requirements, see [“Configuration requirements for shared memory”](#) on page 65.
3. Complete the required preparation tasks. For instructions, see [“Preparing to configure shared memory”](#) on page 72.
4. Create the VIOS logical partitions (hereafter referred to as *paging VIOS partitions*) that you plan to assign to the shared memory pool, and then install the VIOS. For instructions, see [“Creating additional logical partitions”](#) on page 80 and [Installing the VIOS and client logical partitions](#).
5. Create and configure the paging space devices that are owned by the paging VIOS partitions that you plan to assign to the shared memory pool. If you plan to use logical volumes as paging space devices, then create the logical volumes. For instructions, see [“Creating a virtual disk for a VIOS logical partition by using the HMC”](#) on page 122.
6. Verify that the HMC is at version 7 release 3.4.2, or later. For instructions, see [Upgrading your HMC software](#).
7. Ensure that you are a super administrator or operator of the HMC.

About this task

If there is not enough physical memory available in the system to allocate to the shared memory pool, you can release to the hypervisor the physical memory that is currently assigned to logical partitions that use dedicated memory and that are shutdown. The hypervisor can then assign the released physical memory to the shared memory pool.

For more information about managing shared memory pools, see [Managing shared memory pools](#).

What to do next

After you create the shared memory pool, you can create logical partitions that use shared memory. For instructions, see [“Creating additional logical partitions”](#) on page 80.

Creating a logical Host Ethernet Adapter for a running logical partition

If your managed system has a Host Ethernet Adapter (HEA), you can set up a logical partition to use HEA resources by using the Hardware Management Console (HMC) to create a logical Host Ethernet Adapter (LHEA) for the logical partition. A *logical Host Ethernet Adapter (LHEA)* is a representation of a physical HEA on a logical partition. An LHEA allows the logical partition to connect to external networks directly through the HEA. HEAs are also known as Integrated Virtual Ethernet adapters (IVE adapters).

Before you begin

You can add an LHEA dynamically to a running Linux logical partition only if you install Red Hat Enterprise Linux version 5.1, Red Hat Enterprise Linux version 4.6, or a later version of Red Hat Enterprise Linux on the logical partition. To add an LHEA to a Linux logical partition with a distribution other than these distributions, you must shut down the logical partition and reactivate the logical partition using a partition profile that specifies the LHEA.

If a logical partition is not currently running, you can create an LHEA for the logical partition by changing the partition profiles for the logical partition.

About this task

For more information about Logical host Ethernet adapter (LHEA) settings, see [Logical host Ethernet adapter \(LHEA\) settings](#).

Results

When you are done, one or more new Ethernet adapters will be visible to the operating system of the logical partition.

Creating a virtual switch

You can create a virtual switch on a server by using the Hardware Management Console (HMC).

About this task

For more information about creating a virtual switch, see [Adding a virtual network by creating a virtual network bridge](#).

Changing the virtual switch mode setting

When the virtual switch is created, the default setting is the Virtual Ethernet Bridging (VEB) mode. You can change the virtual switch mode to Virtual Ethernet Port Aggregation (VEPA) by using the Hardware Management Console (HMC).

About this task

For more information about changing a virtual switch, see [Changing a virtual switch](#).

Synchronizing the virtual switch mode

When a Virtual I/O Server (VIOS) logical partition is in the shutdown state during the activation of a logical partition or when the external switch is downgraded, the Virtual Station Interface (VSI) profile type information is not updated in the VIOS.

About this task

When any of the VIOS logical partitions that are servicing the virtual switch or when the adjacent connected virtual switches are not in the current switching mode, you must synchronize the switching mode. You can synchronize the switching mode by using the Hardware Management Console (HMC).

You can use the **chhwres** command to synchronize the virtual switch mode.

Creating a Shared Ethernet adapter for a VIOS logical partition by using the HMC

You can create a Shared Ethernet Adapter on the Virtual I/O Server (VIOS) logical partition by using the Hardware Management Console.

About this task

To create a Shared Ethernet Adapter, be sure you meet the following requirements:

- The Hardware Management Console (HMC) must be at version 7 release 3.4.2 or later.
- Ensure the VIOS has one or more physical network devices or Logical Host Ethernet Adapters assigned to the logical partition. If a Logical Host Ethernet Adapter is assigned, the VIOS partition must be configured as the promiscuous logical partition for the Host Ethernet Adapter.
- Ensure a virtual Ethernet adapter is created on the VIOS. For instructions, see [Configuring a virtual Ethernet adapter using the HMC](#).
- If the physical Ethernet adapter that you want to use as the shared adapter has TCP/IP configured, the VIOS must be at version 2.1.1.0 or later. If TCP/IP is not configured, the VIOS can be at any version.
- Ensure that there is a resource monitoring and control connection between the HMC and the VIOS.

Note: If you are using a prior release of the HMC or a prior version of a VIOS (with TCP/IP configured for the virtual Ethernet adapter), see [Configuring virtual Ethernet on the Virtual I/O Server](#) to create a Shared Ethernet Adapter by using the VIOS command-line interface.

For more information about adding a virtual network, see [Managing Virtual Networks](#).

Creating a virtual disk for a VIOS logical partition by using the HMC

You can use the Hardware Management Console (HMC) to create a virtual disk on your managed system. Virtual disks are also known as *logical volumes*.

About this task

To modify virtual storage, be sure you meet the following requirements:

- The HMC must be at version 7.7.4, or later.
- The Virtual I/O Server (VIOS) must be at version 2.2.1.0, or later.
- Ensure that there is a resource monitoring and control connection between the HMC and the VIOS to manage storage.

To create a virtual disk, complete the following steps in the HMC:

Procedure



1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the **PowerVM** area, click **Virtual Storage**. The Virtual Storage page opens with the VIOS partitions listed in a table, in the Virtual Storage Management tab.
4. Select a VIOS and click **Action > Manage Virtual Storage**.
5. In the Virtual Storage Management page, click the **Virtual Disks** tab and click **Create Virtual Disk**.
6. Enter a virtual disk name, select a storage pool or a shared storage pool, and enter the size for the new virtual disk. If you select a shared storage pool, also specify whether you want to use thick or thin storage. By default, the storage type is thin storage. You can optionally assign the disk to a logical partition.
7. Click **OK**

The HMC creates the new virtual disk with your specifications, and the Virtual Disks page is displayed.

Tip: If possible, do not create virtual disks within the *rootvg* storage pool. Create one or more additional storage pools and create the virtual disks using the additional storage pools.

8. Repeat this procedure for each virtual disk that you want to create.
9. To view or change the properties of virtual disks that you created, see [“Changing a virtual disk for a VIOS logical partition by using the HMC”](#) on page 170.

What to do next

These steps are equivalent to using the **mkbdsp** command in the command-line interface.

If there is not enough disk space for the virtual disk, increase the size of the storage pool. For instructions, see [“Changing a storage pool for a VIOS logical partition by using the HMC”](#) on page 172

Creating storage pools

You can use the Hardware Management Console to create a volume-group-based or file-based storage pool on your managed system.

About this task

To create a volume-group-based storage pool, you must assign at least one physical volume to the storage pool. When you assign physical volumes to a storage pool, the Virtual I/O Server erases the information on the physical volumes, divides the physical volumes into physical partitions, and adds the capacity of the physical partitions to the storage pool. Do not add a physical volume to the storage pool if the physical volume contains data that you want to preserve.

To create storage pools, be sure you meet the following requirements:

- The Hardware Management Console must be at version 7 release 3.4.2 or later.
- The Virtual I/O Server must be at version 2.1.1.0 or later.
- Ensure that there is a resource monitoring and control connection between the Hardware Management Console and the Virtual I/O Server.

To create a storage pool, complete the following steps in the Hardware Management Console:

Procedure



1. In the navigation pane, click the **Resources** icon.
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the **PowerVM** area, click **Virtual Storage**. The Virtual Storage page lists the VIOS partitions in the Virtual Storage Management tab.
4. Select a VIOS and click **Action** > **Manage Virtual Storage**.
5. In the Virtual Storage Management page, click the **Storage Pools** tab and click **Create Storage Pool**.
6. Enter a name for the storage pool and select the storage pool type.
7. Enter or select the information required to create the volume-group-based or file-based storage pool, and click **OK** to return to the Storage Pools page.

Note: The new storage pool appears in the table. If you select one or more physical volumes that might belong to a different volume group, the Hardware Management Console displays a warning message to indicate that adding them to the new storage pool can result in data loss. To create the new storage pool with the selected physical volumes, select the Force option, and click **OK**.

Entering the activation code for Active Memory Expansion

You can enable Active Memory Expansion for a server by entering an activation code on the Hardware Management Console (HMC). When you enable Active Memory Expansion, you can configure the logical partitions that run on the server to compress their memory and thus expand their memory capacities.

Before you begin

Before you start, complete the following prerequisite tasks:

1. Complete the required preparation tasks for Active Memory Expansion and ensure that your configuration meets the configuration requirements for Active Memory Expansion. For instructions, see [“Preparing to configure Active Memory Expansion”](#) on page 64.
2. Verify that you have an activation code. You can obtain an activation code from your IBM Sales Representative or from the Capacity on Demand website. To obtain an activation code from the Capacity on Demand website, complete the following steps:
 - a. Go to: <http://www-912.ibm.com/pod/pod>.
 - b. Enter the system type and serial number of the server for which you need the activation code.
 - c. Record the activation code that is displayed on the website.

About this task

To enter the activation code for Active Memory Expansion, complete the following steps by using the HMC:

Procedure



1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system on which you plan to use Active Memory Expansion and click **Actions > View System Properties**. The **Properties** page is displayed.
4. In the **Capacity on Demand** area, click **CoD Functions**.
5. In the Capacity On Demand Functions page, click **Enter CoD Code**.
6. Enter the activation code and click **OK**.
7. From the **Tasks** menu, click **Properties**.
The server Properties window is displayed.
8. Click the **Capabilities** tab.
9. Verify that the **Active Memory Expansion Capable** capability is now set to **True**.
If the capability is set to False, then Active Memory Expansion is not enabled on the server. Obtain a valid activation code to enable Active Memory Expansion on the server.
10. Click **OK**.

What to do next

After you enable Active Memory Expansion on the server, you can configure logical partitions to use Active Memory Expansion. For instructions, see [“Configuring Active Memory Expansion for AIX logical partitions”](#) on page 115.

Reassigning logical partitions to shared processor pools

If you use more than one shared processor pool on your managed system, you can use the Hardware Management Console (HMC) to reassign logical partitions from one shared processor pool to another shared processor pool on your managed system.

Before you begin

You can use this procedure only if the managed system supports more than one shared processor pool and the HMC is at version 7 release 3.2.0, or later.

Any shared processor pool other than the default shared processor pool must be configured before you can assign a logical partition to the shared processor pool. (The default shared processor pool is preconfigured.) For instructions, see [“Configuring shared processor pools”](#) on page 119.

The HMC never allows the sum of the number of reserved processing units for a shared processor pool and the total number of processing units committed to the logical partitions that use the shared processor pool to be greater than the maximum number of processing units for the shared processor pool. (The default shared processor pool has no configured maximum number of processing units. The maximum number of processors available to the default shared processor pool is the total number of active, licensed processors on the managed system minus the number of processors that are assigned to dedicated processor partitions that are set not to share their dedicated processors.)

A shared processor pool cannot contain logical partitions that belong to different workload management groups. You therefore cannot reassign a logical partition with a defined workload management group to a shared processor pool that contains logical partitions that belong to another workload management group. (However, you can reassign a logical partition with a defined workload management group to a shared processor pool that contains only logical partitions that do not have a defined workload management group or that have the same workload management group as the reassigned logical partition.)

About this task

To reassign logical partitions from one shared processor pool to another shared processor pool by using the HMC, follow these steps:

Procedure



1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **Actions** > **View System Properties**. The **Properties** page is displayed.
4. In the **In the PowerVM** area, click **Shared Processor Pool**.
5. In the Shared Processor Pool page, click the **Partitions** tab.
6. Click the name of a logical partition that you want to reassign from one shared processor pool to another shared processor pool.
7. Select the new shared processor pool for the logical partition in the **Pool name (ID)** field and click **OK**.
8. Repeat steps [6](#) and [7](#) for any other logical partitions that you want to reassign from one shared processor pool to another shared processor pool.
9. Click **OK**.

Managing the shared memory pool

By using the Hardware Management Console (HMC), you can change the configuration of the shared memory pool. For example, you can change the amount of physical memory assigned to the shared

memory pool, change the Virtual I/O Server logical partitions that are assigned to the shared memory pool, and add or remove paging space devices to or from the shared memory pool.

Changing the size of the shared memory pool

You can increase or decrease the amount of physical memory assigned to the shared memory pool by using the Hardware Management Console (HMC).

Before you begin

You must be a super administrator or operator of the HMC to change the size of the shared memory pool.

About this task

If there is not enough physical memory available in the system *by which to increase* the amount of memory assigned to the shared memory pool, you can release to the hypervisor the physical memory that is currently assigned to dedicated memory partitions that are shut down. The hypervisor can then assign the released physical memory to the shared memory pool.

If the shared memory pool has insufficient physical memory *by which to decrease* the amount of memory in the shared memory pool, you can release to the hypervisor the I/O entitled memory that is currently assigned to shared memory partitions that are shut down. The hypervisor can then remove the released physical memory from the shared memory pool.

For more information about managing shared memory pools, see [Managing shared memory pools](#).

Adding a paging VIOS partition to the shared memory pool

You can use the Hardware Management Console (HMC) to assign a second Virtual I/O Server (VIOS) logical partition (hereafter referred to as a *paging VIOS partition*) to the shared memory pool.

Before you begin

Before you assign a paging VIOS partition to the shared memory pool, complete the following steps:

1. Verify that only one paging VIOS partition is currently assigned to the shared memory pool.
2. Verify that the paging VIOS partition that is currently assigned to the shared memory pool is running.
3. Verify that the VIOS logical partition that you plan to assign to the shared memory pool is running.
4. Verify that you are a super administrator or an operator of the HMC.

About this task

When you assign a paging VIOS partition to the shared memory pool and both paging VIOS partitions have access to the same paging space devices, those paging space devices become common.

For more information about managing shared memory pools, see [Managing shared memory pools](#).

What to do next

After you assign a second paging VIOS partition to the shared memory pool, complete the following steps:

1. If no common paging space devices are assigned to the shared memory pool, assign them to the shared memory pool. For instructions, see [“Adding and removing paging space devices to and from the shared memory pool”](#) on page 132.
2. Configure the logical partitions that use shared memory to use the paging VIOS partition that you assigned to the shared memory pool. For instructions, see [“Changing the paging VIOS partitions assigned to a shared memory partition”](#) on page 174.

Changing the paging VIOS partitions assigned to the shared memory pool

You can use the Hardware Management Console (HMC) to change the Virtual I/O Server (VIOS) logical partitions (hereafter referred to as *paging VIOS partitions*) that are assigned to the shared memory pool.

Before you begin

Before you change the paging VIOS partitions that are assigned to the shared memory pool, complete the following steps:

1. Shut down all of the shared memory partitions that use the paging VIOS partition that you plan to change. You must shut down all of the shared memory partitions that use the paging VIOS partition (that you plan to change) as the primary paging VIOS partition, and you must shut down all of the shared memory partitions that use the paging VIOS partition (that you plan to change) as the secondary paging VIOS partition. For instructions, see [“Shutting down and restarting logical partitions” on page 137](#).
2. Verify that the VIOS logical partition that you plan to assign to the shared memory pool as a paging VIOS partition is running. (This is the VIOS logical partition to which you plan to change the VIOS assignment of a paging VIOS partition.)
3. Verify that you are a super administrator or an operator of the HMC.

About this task

The following table describes the situations in which you can change a paging VIOS partition.

Table 20. Changing paging VIOS partitions		
State of one paging VIOS partition	State of the other paging VIOS partition	Change options
Running or shut down	None. Only one paging VIOS partition is assigned to the shared memory pool.	You can change the VIOS assignment of the paging VIOS partition. In this situation, you also need to add the paging space devices to which the changed paging VIOS partition has access.
Running	Running	You can change the VIOS assignment of one of the paging VIOS partitions. You cannot change the VIOS assignment of both paging VIOS partitions at the same time.
Running	Shut down	You can change the VIOS assignment of only the paging VIOS partition that is shut down.
Shut down	Running	You can change the VIOS assignment of only the paging VIOS partition that is shut down.

Table 20. Changing paging VIOS partitions (continued)

State of one paging VIOS partition	State of the other paging VIOS partition	Change options
Shut down	Shut down	You cannot change the VIOS assignment of either paging VIOS partition. Instead, you can remove the paging VIOS partition that you do not want to change and then change the VIOS assignment of the remaining paging VIOS partition. In this situation, you also need to add the paging space devices to which the changed paging VIOS partition has access.

When you change the VIOS assignment of a paging VIOS partition, the following configuration changes occur to the paging space devices:

- Paging space devices that were common become independent if only one paging VIOS partition can access them.
- Paging space devices that were common remain common if both paging VIOS partitions can access them. (These are paging space devices to which all three VIOS logical partitions have access. The three VIOS logical partitions are the two VIOS logical partitions that were originally assigned to the shared memory pool as paging VIOS partitions plus the VIOS logical partition that you assigned as a paging VIOS partition when you changed the VIOS assignment of a paging VIOS partition.)
- Paging space devices that were independent become common if both paging VIOS partitions can access them.

When you change the VIOS assignment of a paging VIOS partition, the HMC changes the configuration of the shared memory partitions to use the VIOS logical partition that you assigned as the paging VIOS partition. When you activate the shared memory partition, the HMC automatically reflects the name of the VIOS logical partition that you assigned as the paging VIOS partition in the partition profile. The following examples explain this automatic change in more detail:

- A shared memory partition uses only one paging VIOS partition and you change the VIOS assignment of that paging VIOS partition from VIOS_A to VIOS_B. When you activate the shared memory partition, the HMC automatically shows VIOS_B as the paging VIOS partition in the partition profile.
- Two paging VIOS partitions are assigned to the shared memory pool. VIOS_A is assigned to the shared memory pool as PVP1 and VIOS_B is assigned to the shared memory pool as PVP2. A shared memory partition uses PVP1 as the primary paging VIOS partition and PVP2 as the secondary paging VIOS partition. You change the VIOS assignment of PVP1 from VIOS_A to VIOS_C. When you activate the shared memory partition, the HMC automatically shows VIOS_C as the primary paging VIOS partition and VIOS_B as the secondary paging VIOS partition.

For more information about managing shared memory pools, see [Managing shared memory pools](#).

What to do next

After you change the VIOS assignment of a paging VIOS partition that is assigned to the shared memory pool, complete the following steps:

1. If necessary, assign paging space devices to the shared memory pool. For instructions, see [“Adding and removing paging space devices to and from the shared memory pool” on page 132](#). You might need to add paging space devices in the following situations:
 - You changed the VIOS assignment of the only paging VIOS partition that is assigned to the shared memory pool. The VIOS logical partition that you assigned as the paging VIOS partition has access

to different paging space devices than the VIOS logical partition that was previously assigned as the paging VIOS partition. The paging space devices to which the current paging VIOS partition has access must be assigned to the shared memory pool for the shared memory partitions to use them.

- You removed a paging VIOS partition that was shut down and then changed the VIOS assignment of the other paging VIOS partition that was also shut down. Because you removed a paging VIOS partition from the shared memory pool, you changed the VIOS assignment of the only paging VIOS partition that is assigned to the shared memory pool. The VIOS logical partition that you assigned as the paging VIOS partition has access to different paging space devices than the VIOS logical partition that was previously assigned as the paging VIOS partition. The paging space devices to which the current paging VIOS partition has access must be assigned to the shared memory pool for the shared memory partitions to use them.
 - You changed the VIOS assignment of a paging VIOS partition that provided independent paging space devices to shared memory partitions. The VIOS logical partition that you assigned as the paging VIOS partition has access to different paging space devices than the VIOS logical partition that was previously assigned as the paging VIOS partition. The independent paging space devices to which the current paging VIOS partition has access must be assigned to the shared memory pool for the shared memory partitions to continue to use independent paging space devices.
2. Activate all of the shared memory partitions that you previously shut down so that your changes can take effect. For instructions, see [“Activating a logical partition”](#) on page 133.

Removing a paging VIOS partition from the shared memory pool

You can use the Hardware Management Console (HMC) to remove a Virtual I/O Server (VIOS) logical partition (hereafter referred to as a *paging VIOS partition*) from the shared memory pool.

Before you begin

Before you remove a paging VIOS partition from the shared memory pool, complete the following steps:

1. Verify that two paging VIOS partitions are currently assigned to the shared memory pool.
2. Shut down all of the shared memory partitions that use the paging VIOS partition that you plan to remove. You must shut down all of the shared memory partitions that use the paging VIOS partition (that you plan to remove) as the primary paging VIOS partition, and you must shut down all of the shared memory partitions that use the paging VIOS partition (that you plan to remove) as the secondary paging VIOS partition. For instructions, see [“Shutting down and restarting logical partitions”](#) on page 137.
3. Verify that you are a super administrator or an operator of the HMC.

About this task

The following table describes the situations in which you can remove a paging VIOS partition.

Table 21. Removing paging VIOS partitions		
State of one paging VIOS partition	State of the other paging VIOS partition	Removal options
Running	Running	You can remove either paging VIOS partition.
Running	Shut down	You can remove only the paging VIOS partition that is shut down.
Shut down	Running	You can remove only the paging VIOS partition that is shut down.

Table 21. Removing paging VIOS partitions (continued)

State of one paging VIOS partition	State of the other paging VIOS partition	Removal options
Shut down	Shut down	You can remove either paging VIOS partition; however, you need to reassign the paging space devices to the shared memory pool when you activate the remaining paging VIOS partition. To avoid adding the paging space devices again, you can activate one of the paging VIOS partitions and then remove the other paging VIOS partition.

When you remove a paging VIOS partition from the shared memory pool, the following configuration changes occur:

- Paging space devices that were common become independent.
- The HMC changes the configuration of each shared memory partition to use the remaining paging VIOS partition as the primary and only paging VIOS partition:
 - If a shared memory partition uses only one paging VIOS partition and you remove that paging VIOS partition, the HMC changes the configuration of the shared memory partition to use the remaining paging VIOS partition. When you activate the shared memory partition, the HMC automatically reflects the name of the current paging VIOS partition in the partition profile.

For example, two paging VIOS partitions, VIOS_A and VIOS_B, are assigned to the shared memory pool. A shared memory partition, SMP1, uses only VIOS_A as its paging VIOS partition. You remove VIOS_A from the shared memory pool. When you activate SMP1, the HMC automatically shows VIOS_B as the primary and only paging VIOS partition in the partition profile.

 - If a shared memory partition uses two paging VIOS partitions and you remove a paging VIOS partition, the HMC changes the configuration of the shared memory partition to use the remaining paging VIOS partition as the primary and only paging VIOS partition. When you activate the shared memory partition, the HMC ignores the primary and secondary settings in the partition profile and assigns the remaining paging VIOS partition as the primary and only paging VIOS partition for the shared memory partition. If you want to save the configuration, you can update the partition profile or save the logical partition configuration to a new partition profile.

For more information about managing shared memory pools, see [Managing shared memory pools](#).

What to do next

After you remove a paging VIOS partition from the shared memory pool, complete the following steps:

1. If you removed a paging VIOS partition that was shut down and the other paging VIOS partition was also shut down, complete the following steps:
 - a. Activate the remaining paging VIOS partition. For instructions, see [“Activating a logical partition” on page 133](#).
 - b. Remove the remaining paging space devices from the shared memory pool and assign them again to the shared memory pool. Even though the paging space devices become independent when you remove a paging VIOS partition from the shared memory pool, they cannot be recognized as such until you reassign them to the shared memory pool. For instructions, see [“Adding and removing paging space devices to and from the shared memory pool” on page 132](#).

2. If the paging VIOS partition that you removed was the only paging VIOS partition used by a shared memory partition and the remaining paging VIOS partition does not have access to an available paging space device that meets the size requirements of the shared memory partition, assign such a paging space device to the shared memory pool. For instructions, see [“Adding and removing paging space devices to and from the shared memory pool”](#) on page 132.
3. Activate all of the shared memory partitions that you previously shut down so that your changes can take effect. For instructions, see [“Activating a logical partition”](#) on page 133.

Reinstalling the Virtual I/O Server of a paging VIOS partition

When you reinstall the Virtual I/O Server (VIOS) that is assigned to the shared memory pool (hereafter referred to as a *paging VIOS partition*), you need to reconfigure the shared memory environment. For example, you might need to add the paging space devices again to the shared memory pool.

About this task

The paging VIOS partitions store information about the paging space devices that are assigned to a shared memory pool. The Hardware Management Console (HMC) obtains information about the paging space devices that are assigned to the shared memory pool from the paging VIOS partitions. When you reinstall the VIOS, the information about the paging space devices is lost. For the paging VIOS partitions to regain the information, you must assign the paging space devices again to the share memory pool after you reinstall the VIOS.

The following table shows the reconfiguration tasks that you must perform in the shared memory environment when you resinstall the Virtual I/O Server of a paging VIOS partition.

<i>Table 22. Shared memory reconfiguration tasks for reinstalling the Virtual I/O Server of a paging VIOS partition</i>			
Number of paging VIOS partitions that are assigned to the shared memory pool	Number of paging VIOS partitions for which you want to reinstall the VIOS	Reconfiguration steps	Instructions
1	1	1. Shut down all logical partitions that use shared memory (hereafter referred to as <i>shared memory partitions</i>). 2. Reinstall the VIOS. 3. Add the paging space devices again to the shared memory pool.	1. Shutting down and restarting logical partitions 2. Installing the Virtual I/O Server manually 3. Adding and removing paging space devices to and from the shared memory pool

Table 22. Shared memory reconfiguration tasks for reinstalling the Virtual I/O Server of a paging VIOS partition (continued)

Number of paging VIOS partitions that are assigned to the shared memory pool	Number of paging VIOS partitions for which you want to reinstall the VIOS	Reconfiguration steps	Instructions
2	1	1. Shut down each shared memory partition that uses the paging VIOS partition (that you plan to reinstall) as the primary or secondary paging VIOS partition. 2. Remove the paging VIOS partition from the shared memory pool. 3. Reinstall the VIOS. 4. Add the paging VIOS partition again to the shared memory pool.	1. Shutting down and restarting logical partitions 2. Removing a paging VIOS partition from the shared memory pool 3. Installing the Virtual I/O Server manually 4. Adding a paging VIOS partition to the shared memory pool
2	2	1. Shut down all the shared memory partitions. 2. Reinstall the VIOS of each paging VIOS partition. 3. Add the paging space devices again to the shared memory pool.	1. Shutting down and restarting logical partitions 2. Installing the Virtual I/O Server manually 3. Adding and removing paging space devices to and from the shared memory pool

Adding and removing paging space devices to and from the shared memory pool

After you create the shared memory pool, you can add and remove paging space devices to and from the shared memory pool by using the Hardware Management Console (HMC).

Before you begin

Before you add a paging space device, complete the following tasks:

1. Configure the paging space device to the Virtual I/O Server (VIOS) logical partitions (hereafter referred to as *paging VIOS partitions*) that are assigned to the shared memory pool. If you plan to use logical volumes as paging space devices, then create the logical volumes. For instructions, see [“Creating a virtual disk for a VIOS logical partition by using the HMC” on page 122](#).
2. Verify that all paging VIOS partitions are running.

Before you remove a paging space device, complete the following tasks:

- If no logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) is using the paging space device, verify that the paging space device is inactive.

- If a shared memory partition is using the paging space device, verify that the shared memory partition is shut down.
- Verify that all paging VIOS partitions are running.

You must be a super administrator or operator of the HMC to add and remove paging space devices to and from the shared memory pool.

About this task

For more information about managing shared memory pools, see [Managing shared memory pools](#).

Deleting the shared memory pool

If you no longer want any of the logical partitions to use shared memory, you can delete the shared memory pool by using the Hardware Management Console (HMC).

Before you begin

Before you start, remove all of the logical partitions that use shared memory (hereafter referred to as *shared memory partitions*) from the shared memory pool by completing one of the following tasks:

- Delete all of the shared memory partitions. For instructions, see [“Deleting a logical partition” on page 114](#).
- Change all of the shared memory partitions to dedicated memory partitions. For instructions, see [“Changing the memory mode of a logical partition” on page 176](#).

About this task

For more information about managing shared memory pools, see [Managing shared memory pools](#).

Managing logical partitions

You can manage the configuration of your logical partitions by using the Hardware Management Console (HMC). The HMC allows you to adjust the hardware resources that are used by each logical partition.

Activating a logical partition

You must activate a logical partition before you can use the logical partition. When you use the Hardware Management Console (HMC), you can activate a logical partition based on its current configuration or you can activate a logical partition by activating a partition profile.

About this task

For more information about activating a logical partition, see [Activating partitions](#).

Activating a partition profile

You can activate a partition profile by using the Hardware Management Console (HMC). When you activate a partition profile, you activate a logical partition. The system commits resources to the logical partition based on the configuration in the partition profile and starts the operating system or software that is installed on the logical partition.

Before you begin

When you activate a logical partition by activating a partition profile, you must select a partition profile. A *partition profile* is a record on the HMC that specifies a possible configuration for a logical partition.

If you plan to activate a logical partition that uses virtual resources provided by the Virtual I/O Server, you must first activate the Virtual I/O Server (VIOS) logical partition that provides the virtual resources.

If you plan to activate a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*), you must first activate at least one VIOS logical partition that meets the following criteria:

- The VIOS logical partition (hereafter referred to as a *paging VIOS partition*) must provide access to an available paging space device that meets the size requirements of the shared memory partition.
- The paging VIOS partition must be assigned to the shared memory pool.

If the shared memory partition is configured with redundant paging VIOS partitions, activate both paging VIOS partitions before you activate the shared memory partition.

When you activate a shared memory partition and the shared memory pool does not contain enough physical memory required for activation, you can release to the hypervisor the physical memory that is currently assigned to other shared memory partitions that are shut down. The hypervisor can then assign the released physical memory to the shared memory partition that you want to activate.

When the partition profile contains a cable card, the partition activation fails. You must remove the cable card from the profile before you activate the logical partition, because a slot that contains a cable card cannot be partitioned.

About this task

To activate a partition profile by using the HMC, follow these steps:

Procedure

1. In the navigation pane, open **Systems Management > Servers**, and click the system on which the logical partition is located.
2. In the work pane, select the logical partition that you want to activate.
3. From the Tasks menu, click **Operations > Activate > Profile**.
4. If you want to install the VIOS software as part of the activation process of a logical partition, complete the following steps:
 - a) Click **Yes** as the value for the **Install Virtual I/O Server as part of activation process** field.
 - b) Select the partition profile that you want to use to activate the logical partition.
 - c) Click **OK**.

The **Discovering Network Adapters** window is displayed because it might take some time to load the network adapters.
 - d) On the **Install Virtual I/O Server** page, select the VIOS installation source, and complete the required fields.
 - e) Click **OK**.

The installation progress pane displays the status of the VIOS installation in the progress bar. To view the details about the progress of the installation, click the **Details** tab.
 - f) Click **Close**.

A message is displayed that the VIOS installation was successful. If you selected **NIM Server** as the installation source, the NIM installation starts after you click **Close** in the installation progress pane. To view the progress of the NIM installation from a virtual terminal, click **Popup Console**. When the NIM installation is complete, a message is displayed that the installation was successful.
 - g) Click **OK**.
5. If you want the HMC to open a terminal window or console session for the logical partition when the logical partition is activated, click **Open a terminal window or console session**.

Note: This option is disabled when you select **Yes** as the value for the **Install Virtual I/O Server as part of activation process** field.

6. If you want to use a keylock position, boot mode, or paging VIOS redundancy configuration that is different from the keylock position, boot mode, or paging VIOS redundancy configuration that is specified in the partition profile, complete the following steps:
 - a) Click **Advanced**.
 - b) Select the desired keylock position, boot mode, or paging VIOS redundancy configuration.
 - c) Click **OK**.

7. Click **OK**.

If the logical partition that you want to activate is a shared memory partition and there is not enough physical memory in the shared memory pool by which to activate the shared memory partition, the Release Memory Resources window is displayed.

8. Select shared memory partitions that are shut down until the available memory is equal to or greater than the requested memory and click **OK**.

Activating a logical partition based on its current configuration

You can use the Hardware Management Console (HMC) to activate a logical partition based on its current configuration instead of a partition profile. When you activate the logical partition, the system commits resources to the logical partition based on the current configuration of the logical partition and starts the operating system or software that is installed on the logical partition. Logical partitions start faster when activated based on their current configuration data than when activated with a partition profile.

Before you begin

You cannot activate a logical partition based on its current configuration if one of the following conditions is true:

- The state of the logical partition is such that the logical partition is not capable of starting. To activate the logical partition based on its current configuration, change the state of the logical partition such that it is capable of starting.
- There is no active partition profile associated with the logical partition. For example, a newly created logical partition that has never been activated does not have an active partition profile. This logical partition cannot be activated based on its current configuration because its current configuration has no resources. The first time you activate a logical partition, you must activate it by activating a partition profile.

If you plan to activate a logical partition that uses virtual resources provided by the Virtual I/O Server, you must first activate the Virtual I/O Server (VIOS) logical partition that provides the virtual resources.

If you plan to activate a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*), you must first activate at least one VIOS logical partition that meets the following criteria:

- The VIOS logical partition (hereafter referred to as a *paging VIOS partition*) must provide access to an available paging space device that meets the size requirements of the shared memory partition.
- The paging VIOS partition must be assigned to the shared memory pool.

If the shared memory partition is configured with redundant paging VIOS partitions, activate both paging VIOS partitions before you activate the shared memory partition.

When you activate a shared memory partition and the shared memory pool does not contain enough physical memory required for activation, you can release to the hypervisor the physical memory that is currently assigned to other shared memory partitions that are shut down. The hypervisor can then assign the released physical memory to the shared memory partition that you want to activate.

On a HMC that is at a version earlier to Version 7.8.0, if the resource configuration field of the partition is set to **Not configured**, activating a logical partition with current configuration results in an error. On a HMC

at Version 7.8.0, or later, if the resource configuration field is set to **Not configured**, and the partition has a last valid configuration profile, then that profile is used to activate the partition.

About this task

To activate a logical partition based on its current configuration by using the HMC, complete the following steps:

Procedure



1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions > View Partition Properties**.
4. Click the **Partition Actions > Operations > Activate**.
5. In the Activate page, select **Current Configuration** as the value for the **Partition Configuration** field.
6. Click **Finish**.

Viewing the resource configuration status of a logical partition

You can view the resource configuration status of a logical partition by using the Hardware Management Console (HMC).

About this task

To view the resource configuration of a logical partition by using the HMC, follow these steps:

Procedure



1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions > View Partition Properties**.
4. Click the **General** tab
When the **Resource Configuration** field displays **Configured**, the partition can be activated by using the current configuration profile. When the **Resource Configuration** field displays **Not Configured**, the partition is activated the by using the last valid configuration that was stored as a profile.
5. Click **OK**.

Applying a profile to a logical partition

On a Hardware Management Console (HMC) that is at Version 7 Release 7.8.0, or later, you can apply a profile to a logical partition without powering on the logical partition by using the HMC command-line interface.

Procedure

From the HMC command line, type the following command:

```
chsyscfg -r lpar -m managed system -o apply -n profile name
```

Where:

- *managed system* is the name of the server on which the logical partition is located.
- *profile name* is the name of the partition profile that applied to the logical partition.

Activating a system profile

You can activate many logical partitions at a time by using the Hardware Management Console (HMC) to activate a system profile. A *system profile* is an ordered list of partition profiles. When you activate a system profile, the managed system attempts to activate the partition profiles in the system profile in the order in which the partition profiles are listed.

About this task

Restriction: You cannot activate a system profile that contains partition profiles that specify shared memory.

To activate a system profile using the HMC, follow these steps:

Procedure



1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **System Actions > Legacy > Manage System Profiles**.
4. In the Manage System Profiles page, select the profile from the list and click **Actions > Activate**.
5. Select the desired activation settings for the system profile and click **Continue**.

Shutting down and restarting logical partitions

You can shut down and restart logical partitions running on systems that are managed by a Hardware Management Console (HMC).

Shutting down and restarting AIX in a logical partition

You can shut down and restart AIX in a logical partition using the Hardware Management Console (HMC).

Shutting down AIX logical partitions

You can shut down AIX logical partitions using the Hardware Management Console (HMC).

About this task

To shut down an AIX logical partition, complete the following steps from the HMC:

Procedure



1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
4. In the work pane, select the partition and click **Actions > Shutdown**.
5. Select one of the following options:

Option	Description
Operating System	The HMC issues the AIX shutdown command to shut down the logical partition. This option is available only when the operating system is running, and not when the logical partition is in an Open Firmware state.
Operating System Immediate	The HMC issues the AIX shutdown -F command to shut down the logical partition as quickly as possible, bypassing messages to other users. This option

Option	Description
	is available only when the operating system is running, and not when the logical partition is in an Open Firmware state.
Delayed	The logical partition waits a predetermined amount of time to shut down. This allows the logical partition time to end jobs and write data to disks.
Immediate	The logical partition shuts down without any preset delay.

6. Click **OK**.

Restarting AIX logical partitions

You can restart AIX logical partitions using the Hardware Management Console (HMC). Restarting a logical partition shuts the logical partition down and then starts it again.

About this task

To restart an AIX logical partition, complete the following steps from the HMC:

Procedure



1. In the navigation pane, click the **Resources** icon.
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions > Restart**.
4. Select one of the following options:

Option	Description
Operating System	The HMC issues the AIX shutdown -r command to shut down and restart the logical partition. This option is available only when the operating system is running, and not when the logical partition is in an Open Firmware state.
Operating System Immediate	The HMC issues the AIX shutdown -F -r command to shut down and restart the AIX logical partition as quickly as possible, bypassing messages to other users. This option is available only when the operating system is running, and not when the logical partition is in an Open Firmware state.
Immediate	The logical partition is restarted as quickly as possible, without notifying the logical partition.
Dump	The HMC initiates a main storage or system memory dump on the logical partition and restarts the logical partition after the dump.

5. Click **OK**.

Shutting down IBM i logical partitions

The correct way to shut down an IBM i logical partition safely is from an IBM i command line.

If you cannot shut down the IBM i logical partition from an IBM i command line, you can shut down the IBM i logical partition from the Shut Down Partition window on your HMC or from the remote control panel on the Operations Console. Using these methods can cause an abnormal shutdown and can result in loss of data.

Before you shut down an IBM i logical partition, you must perform all of the basic IBM i shutdown tasks. For example, all other users must be signed off of the IBM i logical partition before you can shut it down. If you shut down the IBM i logical partition without completing all of the required tasks, you can cause damage to data or cause the system to behave in unpredictable ways. For instructions, see [Basic system operations](#).

Shutting down IBM i logical partitions by using the HMC

You can shut down IBM i logical partitions by using the Hardware Management Console (HMC).

Before you begin

Before you shut down the IBM i logical partition, complete the following tasks:

1. If an integrated server is active on the system, shut down the integrated server using IBM i options.
2. Ensure that all jobs are completed and all applications are ended.
3. Ensure that your partition profiles are updated with any dynamic partitioning resource changes that you want to keep when you restart the logical partition.

About this task

The correct way to shut down an IBM i logical partition from the HMC is to open an HMC 5250 emulator session and run the Power Down System (PWRDWNSYS) command.

To shut down an IBM i logical partition from the HMC, follow these steps:

Procedure



1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. Select the logical partition and click **Actions > Console > Open Dedicated 5250 Console**.
4. From the IBM i command line in the emulator session, type PWRDWNSYS OPTION (*CNTRL) DELAY (600) and press Enter.

The system will only shut down the IBM i logical partition you selected. The PWRDWNSYS command does not affect other IBM i logical partitions on your system. If you enter the PWRDWNSYS command with the RESTART(*YES) option, the operating system restarts, and the resource specifications of the logical partition remain the same. If you do not use the RESTART(*YES) option, then the logical partition shuts down completely, and other logical partitions will be able to take and use the resources that were used by the logical partition. Also, when you reactivate the logical partition using a partition profile, the partition profile overlays the resource specifications of the logical partition with the resource specifications in the partition profile. Any resource changes that you made to the logical partition using dynamic partitioning are lost when you reactivate the logical partition using a partition profile. If the logical partition is set to start automatically when the managed system starts, you can preserve the resource specifications on that logical partition by restarting the entire managed system using the Partition autostart power-on mode. When the logical partitions start automatically, the logical partitions have the resource specifications that the logical partitions had when you shut down the managed system.

5. If the PWRDWNSYS command does not work, you can use either a delayed shutdown or an immediate shutdown to shut down the IBM i logical partition.



Attention: Using these methods can cause an abnormal shutdown and can result in loss of data.

Performing a delayed shutdown of an IBM i logical partition

You can perform a delayed shutdown of a logical partition using the Hardware Management Console (HMC). Using delayed shutdown is equivalent to using the power button on the remote control panel. Use

delayed shutdown only when you must shut down a logical partition, and the PWRDWNSYS command does not work.

Before you begin

When you use the delayed shutdown option, the logical partition waits a predetermined amount of time to shut down. This allows the logical partition time to end jobs and write data to disks. If the logical partition is unable to shut down within the predetermined amount of time, it will end abnormally and the next restart might take a long time.

About this task

To perform a delayed shutdown of an IBM i logical partition using the HMC, complete the following steps:

Procedure



1. In the navigation pane, click the **Resources** icon.
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions > Shutdown**.
4. In the Shut Down Partitions page, select **Delayed** and click **OK**.

Performing an immediate shutdown of an IBM i logical partition

When you perform the immediate shutdown option by using the Hardware Management Console (HMC), the system shuts down without any preset delay. Using immediate shutdown is equivalent to using function 8 on the remote control panel.

Before you begin



Attention: Using immediate shutdown can cause an abnormal IPL of the IBM i logical partition and possibly cause loss of data. Use immediate shutdown only when an IBM i logical partition cannot shut down using PWRDWNSYS or delayed shutdown.

About this task

To perform an immediate shutdown of an IBM i logical partition using the HMC, complete the following steps:

Procedure



1. In the navigation pane, click the **Resources** icon.
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions > Shutdown**.
4. In the Shut Down Partitions page, select **Immediate** and click **OK**.

Shutting down IBM i logical partitions by using Operations Console

You can shut down IBM i logical partitions by using Operations Console.

Before you shut down the IBM i logical partition, complete the following:

1. If an integrated server is active on the system, shut down the integrated server by using IBM i options.
2. Ensure that all jobs are completed and all applications are ended.

3. Ensure that your partition profiles are updated with any dynamic partitioning resource changes that you want to keep when you restart the logical partition.

The correct way to shut down a logical partition is by using the control language (CL) command Power Down System (PWRDWNSYS).

From an IBM i command line, type PWRDWNSYS OPTION (*CNTRLD) DELAY (600) and press Enter. The system will only shut down the IBM i logical partition you selected. The PWRDWNSYS command does not affect other IBM i logical partitions on your system.

If you enter the PWRDWNSYS command with the RESTART(*YES) option, the operating system restarts, and the resource specifications of the logical partition remain the same. If you do not use the RESTART(*YES) option, then the logical partition shuts down completely, and other logical partitions will be able to take and use the resources that were used by the logical partition. Also, when you reactivate the logical partition using a partition profile, the partition profile overlays the resource specifications of the logical partition with the resource specifications in the partition profile. Any resource changes that you made to the logical partition using dynamic partitioning are lost when you reactivate the logical partition using a partition profile. If the logical partition is set to start automatically when the managed system starts, you can preserve the resource specifications on that logical partition by restarting the entire managed system using the Partition autostart power-on mode. When the logical partitions start automatically, the logical partitions have the resource specifications that the logical partitions had when you shut down the managed system.

If the PWRDWNSYS command does not work, you can use the remote control panel through Operations Console to use control panel functions through a PC. The graphical user interface of the remote control panel looks similar to the physical control panel. The remote control panel installs through Operations Console. Using the remote control panel to shut down the IBM i logical partition can result in an abnormal IPL and loss of data.

Delayed shutdown

Use delayed shutdown only when you must shut down a logical partition, and when the PWRDWNSYS command does not work.

When you use the delayed shutdown option, the logical partition waits a predetermined amount of time to shut down. This allows the logical partition time to end jobs and write data to disks. If the logical partition is unable to shut down within the predetermined amount of time, it will end abnormally and the next restart might take a long time.

Immediate shutdown

Use immediate shutdown only when an IBM i logical partition cannot shut down using PWRDWNSYS or delayed shutdown.

When you use the immediate shutdown option, the system turns off without any preset delay.



Attention: This might cause an abnormal IPL of the IBM i logical partition and possibly cause loss of data.

Use the remote control panel to perform a delayed shutdown or an immediate shutdown. The power button will start a delayed shutdown and function 8 will start an immediate shutdown of a system.

Restarting and shutting down IBM i in a logical partition

At times you will need to perform an initial program load (IPL) or shut down an IBM i logical partition. For example, if you want to apply a delayed fix to IBM i, you must perform an IPL before IBM i can apply the fix.

The preferred method for restarting and shutting down IBM i logical partitions is through the IBM i command line. The Hardware Management Console (HMC) does not shut down the IBM i operating system before it shuts down the logical partition. Using the HMC to restart or shut down an IBM i logical partition can result in an abnormal IPL and the loss of data. However, you might need to use the HMC to

change the operating mode or IPL type of the IBM i logical partition before you restart or shut down the IBM i logical partition using the IBM i command line.

It is important to remember that, when you perform an IPL of an IBM i logical partition, you are powering off only the logical partition and not the entire managed system. Other logical partitions on your managed system continue to run when you perform an IPL on the IBM i logical partition. However, when you shut down the last logical partition that is running on a managed system, then the managed system is set to power off automatically by default. If you want, you can set the managed system properties on the HMC so that the managed system remains powered on when you shut down the last running logical partition.

Changing the operating mode for an IBM i logical partition

You can change the operating mode for an IBM i logical partition using the Hardware Management Console (HMC). The operating mode for an IBM i logical partition determines the number of options that are presented to the operator for consideration during and after the initial program load (IPL). It can also secure (lock) the control panel to prevent an unauthorized or inadvertent IPL from the control panel.

About this task

To change the IBM i operating mode of a logical partition by using the HMC, follow these steps:

Procedure



1. In the navigation pane, click the **Resources** icon.
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions > View Partition Properties**.
4. Click the **General** tab and set **Key Lock position** to your preference, and click **OK**.

Changing the IPL type for an IBM i logical partition

When you use the Hardware Management Console (HMC) to change the IPL type, the managed system loads the Licensed Internal Code and IBM i from the location specified by the IPL type. The IPL type is also known as the IPL source, because each IPL type is associated with a different IPL source.

Before you begin

You can choose a separate IPL type for each IBM i logical partition.



Attention: Only use IPL type C under the direction of your service representative. Severe data loss can occur with incorrect use of this function.

About this task

To change the IBM i IPL type of a logical partition using the HMC, follow these steps:

Procedure



1. In the navigation pane, click the **Resources** icon.
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions > View Partition Properties**.
4. Click the **General** tab and set **IPL Source** to your preference, and click **OK**.

Shutting down and restarting Linux in a logical partition

You can shut down and restart Linux logical partitions or the Linux operating system by using the Hardware Management Console (HMC).

Shutting down Linux logical partitions

You can shut down Linux logical partitions and the Linux operating system using the Hardware Management Console (HMC).

About this task

To shut down a Linux logical partition, complete the following steps from the HMC:

Procedure



1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
4. In the work pane, select the partition and click **Actions > Shutdown**.
5. Select one of the following options:

Option	Description
Operating System	The HMC issues the Linux shutdown -h +1 command to shut down the logical partition. This option is available only when the operating system is running, and not when the logical partition is in an Open Firmware state.
Operating System Immediate	The HMC issues the Linux shutdown -h now command to shut down the logical partition as quickly as possible, bypassing messages to other users. This option is available only when the operating system is running, and not when the logical partition is in an Open Firmware state.
Delayed	The logical partition waits a predetermined amount of time to shut down. This allows the logical partition time to end jobs and write data to disks.
Immediate	The logical partition shuts down without any preset delay.

6. Click **OK**.

Restarting Linux logical partitions

You can restart Linux logical partitions or the Linux operating system using the Hardware Management Console (HMC). Restarting a logical partition shuts the logical partition down and then starts it again.

About this task

To restart a Linux logical partition, complete the following steps from the HMC:

Procedure



1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions > Restart**.
4. Select one of the following options:

Option	Description
Operating System	The HMC issues the Linux shutdown -r +1 command to shut down and restart the logical partition. This option is available only when the operating system is running, and not when the logical partition is in an Open Firmware state.
Operating System Immediate	The HMC issues the Linux shutdown -r now command to shut down and restart the logical partition as quickly as possible, bypassing messages to other users.
Immediate	The logical partition is restarted as quickly as possible, without notifying the logical partition.
Dump	The HMC allows the Linux operating system on the Linux logical partition to run a diagnostic procedure. After the diagnostic procedure is complete, the logical partition restarts. The exact diagnostic procedure depends upon which Linux operating system is installed on the logical partition and how the operating system is set. The operating system might run an OS debugger, the operating system might perform a main storage or system memory dump on the logical partition, or the operating system might not be set to run any diagnostic procedure at all.

5. Click **OK**.

Shutting down and restarting Virtual I/O Server in a logical partition

You can shut down and restart Virtual I/O Server by using the Hardware Management Console (HMC).

Shutting down Virtual I/O Server logical partitions by using the HMC

You can shut down Virtual I/O Server logical partitions by using the Hardware Management Console (HMC). You can shut down the Virtual I/O Server immediately or delay the shutdown.

Before you begin

Before you shut down the Virtual I/O Server logical partition, complete the following tasks:

- If the client logical partitions that use storage and networking virtual resources provided by the Virtual I/O Server are not configured to use virtual resources provided by a redundant Virtual I/O Server, then shut down the client logical partitions.
- Shut down each shared memory partition that accesses its paging space device using only the Virtual I/O Server logical partition that you plan to shut down. If you shut down the Virtual I/O Server (VIOS) logical partition (hereafter referred to as a *paging VIOS partition*) before you shut down the shared memory partitions and a shared memory partition attempts to access memory that is located on its paging space device, the shared memory partition might fail.

If a shared memory partition accesses its paging space device redundantly through two paging VIOS partitions, you do not need to shut down the shared memory partition. When you shut down the paging VIOS partition, the shared memory partition accesses its paging space device through the other paging VIOS partition.

About this task

To shut down a Virtual I/O Server logical partition, complete the following steps from the HMC:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.

3. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
4. In the work pane, select the partition and click **Actions > Shutdown**.
5. Select one of the following options:

Option	Description
Operating System	The HMC issues the Virtual I/O Server shutdown command to shut down the logical partition. This option is available only when the operating system is running, and not when the logical partition is in an Open Firmware state.
Operating System Immediate	The HMC issues the Virtual I/O Server shutdown -force command to shut down the logical partition as quickly as possible, bypassing messages to other users. This option is available only when the operating system is running, and not when the logical partition is in an Open Firmware state.
Delayed	The logical partition waits a predetermined amount of time to shut down. This allows the logical partition time to end jobs and write data to disks.
Immediate	The logical partition shuts down without any preset delay.

6. Click **OK**.

Restarting Virtual I/O Server logical partitions by using the HMC

You can restart Virtual I/O Server logical partitions by using the Hardware Management Console (HMC). Restarting a Virtual I/O Server logical partition shuts down the Virtual I/O Server logical partition and then starts it again.

Before you begin

Before you shut down the Virtual I/O Server logical partition, complete the following tasks:

- If the client logical partitions that use storage and networking virtual resources provided by the Virtual I/O Server are not configured to use virtual resources provided by a redundant Virtual I/O Server, then shut down the client logical partitions.
- Shut down each shared memory partition that accesses its paging space device using only the Virtual I/O Server logical partition that you plan to shut down. If you shut down the Virtual I/O Server (VIOS) logical partition (hereafter referred to as a *paging VIOS partition*) before you shut down the shared memory partitions and a shared memory partition attempts to access memory that is located on its paging space device, the shared memory partition might fail.

If a shared memory partition accesses its paging space device redundantly through two paging VIOS partitions, you do not need to shut down the shared memory partition. When you shut down the paging VIOS partition, the shared memory partition accesses its paging space device through the other paging VIOS partition.

About this task

To restart a Virtual I/O Server logical partition, complete the following steps from the HMC:

Procedure



1. In the navigation pane, click the **Resources** icon.
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions > Restart**.
4. Select one of the following options:

Option	Description
Operating System	The HMC issues the Virtual I/O Server shutdown -restart command to shut down and restart the logical partition. This option is available only when the operating system is running, and not when the logical partition is in an Open Firmware state.
Operating System Immediate	The HMC issues the Virtual I/O Server shutdown -force -restart command to shut down and restart the logical partition as quickly as possible, bypassing messages to other users. This option is available only when the operating system is running, and not when the logical partition is in an Open Firmware state.
Immediate	The logical partition is restarted as quickly as possible, without notifying the logical partition.
Dump	The HMC initiates a main storage or system memory dump on the logical partition and restarts the logical partition after the dump.

5. Click **OK**.

Results

After the Virtual I/O Server restarts, complete the following tasks:

- Activate the client logical partitions that use storage and networking virtual resources provided to them by the Virtual I/O Server.
- Activate each shared memory partition that accesses its paging space device by using only the paging VIOS partition that you restarted.

Partition time power-on

If the operating system on a logical partition is scheduled to start at a specified time, the server starts automatically if the server is not already powered on. The hosting partition on a Virtual Partition Manager managed system also starts automatically if the partitions are not already running. However, other logical partitions are not started automatically at the specified time, even if the logical partitions are set to auto start at system power-on.

Procedures for scheduling operating system power-on

Each operating system has its own procedure for scheduled power-on. For information about how to schedule power-on for a specific operating system, see the link for that operating system in the following table.

<i>Table 23. Procedures for scheduling operating system power-on</i>	
Operating system	Procedure for scheduling operating system power-on
AIX	Run the shutdown command, and specify the time at which the restart must occur, by using the -t flag. For more information, see the shutdown command page .
IBM i	Scheduling a system shutdown and restart
Linux	Set power-on time

Hardware behavior for different system configurations

When an operating system starts at a scheduled time, the configuration of the managed system determines how the operating system starts and what is started with the operating system. The following table shows the hardware behavior for each managed system configuration.

Table 24. Hardware behavior when an operating system is scheduled to power on

System configuration	Hardware behavior when an operating system is scheduled to power on
Manufacturing default configuration (MDC), where a single partition owns all the resources of the system	The following activities occur on the server that contains the operating system that is scheduled to power on: 1. The server powers on. 2. The time power-on MDC partition starts.
Server is managed by Virtual Partition Manager	The following activities occur on the server that contains the client partition that is scheduled to power on: 1. The server powers on if the server is not already on, and the hosting partition starts if the hosting partition has not already started. 2. The time power-on partition or partitions starts. Partitions that are defined to auto start with the server power will not power on automatically when the partition that is scheduled to power on starts. To set multiple partitions to start at the same time, you must set the operating system of each partition to start then.
Server is managed by the Hardware Management Console (HMC)	The following activities occur on the server that contains the partition that is scheduled to power on: 1. The server powers on, if the server is not already on. 2. The time power-on partition or partitions starts. Partitions that are defined to auto start with the server power will not power on automatically when the partition that is scheduled to power on starts. To set multiple partitions to start at the same time, you must set the operating system of each partition to start then.

Managing partition profiles for logical partitions

You can manage the partition profiles for your logical partitions using the Hardware Management Console (HMC). You can change the resource specifications stored in your partition profiles as your needs change.

Copying a partition profile

You can create a copy of an existing partition profile using the Hardware Management Console (HMC). After you create a copy of the existing partition profile, you can change the resource allocations within the new partition profile. This allows you to create multiple, nearly identical partition profiles without having to re-enter all of the resource allocations repeatedly.

About this task

To copy a partition profile using the HMC, follow these steps:

Procedure



1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions > Profiles > Manage Profiles**.
4. Select the partition profile that you want to copy and click **Actions > Copy**.
5. Enter the name of the new partition profile into **New profile name** and click **OK**.

Results

Viewing the vNIC properties in a partition profile

You can view the properties of a virtual Network Interface Controller (vNIC) in a partition profile by using the Hardware Management Console (HMC). You can only view the properties of the vNIC and you cannot create, edit, or remove a vNIC in a partition profile.

About this task

A virtual Network Interface Controller (vNIC) is a type of virtual Ethernet adapter that can be configured on client logical partitions. Each vNIC is backed by a single root I/O virtualization (SR-IOV) logical port that is owned by the VIOS. When the HMC is at version 8.6.0, or later, the firmware is at level FW860, or later, and the VIOS is at version 2.2.5.0, or later, a dedicated vNIC can have multiple SR-IOV logical ports on different physical ports as backing devices, and they can be hosted by the same or different Virtual I/O Servers.

To view the vNIC properties in a partition profile by using the HMC, complete the following steps:

Procedure



1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions > Profiles > Manage Profiles**.
4. Click the partition profile name that you want to view.
5. In the Logical Partition Profile Properties page, click the **Virtual Adapters** tab. If vNICs exist in the profile, they will be displayed. Click the Adapter ID of the vNIC adapter for which you want to view the properties.
In the **General** tab of the Virtual NIC Adapter Properties page, you can view the **Adapter ID**, **Port ID**, **Hosting VIOS**, **Failover Priority**, and **Capacity** fields. The **Backing Devices** table displays all backing devices that are available. When there is more than one backing device, the **Auto Priority Failback** field is displayed.
6. In the Virtual NIC Adapter Properties page, click the **Advanced** tab.
You can view the **Port VLAN ID**, **PVID Priority**, **VLAN Restrictions**, **MAC Address**, and **MAC Address Restrictions** fields.

Changing partition profile properties

You can change the properties of a partition profile using the Hardware Management Console (HMC). Changing the properties of a partition profile changes the resource amounts that are assigned to a logical partition when you shut down and restart the logical partition using the changed partition profile.

Before you begin

A partition profile stores the required number of processors, memory, and hardware resources assigned to that profile. Any partition profile property changes are not applied to the logical partition until you activate the partition profile.

If you plan to change a partition profile that specifies dedicated memory to a partition profile that specifies shared memory, be aware of the following actions:

- The HMC automatically deletes all of the physical I/O adapters specified in the partition profile. You can assign only virtual adapters to logical partitions that use shared memory.
- You must specify shared processors. Logical partitions that use shared memory must also use shared processors.

About this task

To change partition profile properties using the HMC, follow these steps:

Procedure



1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions > Profiles > Manage Profiles**.
4. Select the partition profile that you want to change and click **Actions > Edit**.

To add, remove, or change the vNIC adapter settings, you can run the **chsyscfg** command from the HMC command line. To add vNIC backing devices to a partition or to remove vNIC backing devices from a partition, and to change the vNIC auto-failback policy or to change the vNIC backing device failover policy, run the **chhwres** command from the HMC command line.

When the HMC is at Version 9.1.0, or later, you can use the *max_capacity* field in the vNIC backing device attribute of the **chsyscfg** command to configure vNIC backing devices. You can also use the *max_capacity* attribute of the **chsyscfg** command to configure a single root I/O virtualization (SR-IOV) Ethernet logical port.

5. Make the appropriate changes and click **OK**.

What to do next

If you created at least one virtual fibre channel adapter, complete the following tasks to connect the logical partition to its storage:

1. Activate the logical partition. When you activate the logical partition, the HMC assigns a pair of worldwide port names (WWPNs) to the virtual fibre channel adapter. For instructions, see [“Activating a logical partition” on page 133](#).
2. Restart the Virtual I/O Server (that provides the connection to a physical fibre channel adapter) or run the **syscfg** command. This enables the Virtual I/O Server to recognize the WWPNs of the virtual fibre channel adapter on the client logical partition. For instructions, see [“Restarting Virtual I/O Server logical partitions by using the HMC” on page 145](#).
3. Assign the virtual fibre channel adapter on the client logical partition to a physical port of a physical fibre channel adapter. For instructions, see [“Changing virtual Fibre Channel for a Virtual I/O Server by using the HMC” on page 173](#).

Related information

[chsyscfg command](#)

[chhwres command](#)

[lshwres command](#)

Deleting a partition profile

You can delete a partition profile using the HMC Hardware Management Console (HMC). This allows you to remove partition profiles that you no longer use.

Before you begin

Note: You cannot delete a partition profile that is the default partition profile for the logical partition. If the partition profile you want to delete is the default partition profile, you must first change the default profile to another partition profile.

About this task

To delete a partition profile using the HMC, follow these steps:

Procedure



1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions > Profiles > Manage Profiles**.
4. Select the partition profile that you want to delete and click **Actions > Delete**.
5. Click **OK** to confirm.

Adding the PPT ratio to the partition profile

When the Hardware Management Console (HMC) is at Version 9.1.0, or later, you can add the Physical Page Table (PPT) ratio to the partition profile.

About this task

You must be a super administrator to complete this task and the server must be in the operating state.

You can run the **chsyscfg** command from the HMC command-line interface to add the PPT ratio to the logical partition profile. To view the PPT ratio of a logical partition profile, run the **lssyscfg** command.

Related information

[lssyscfg command](#)

[chsyscfg command](#)

Managing system profiles

You can manage the system profiles on your managed system using the Hardware Management Console (HMC). You can change the logical partitions and partition profiles specified within the system profiles as the logical partitions change on your managed system.

Copying a system profile

You can use the Hardware Management Console (HMC) to create a copy of an existing system profile. After you create a copy of the existing system profile, you can change the partition profiles that are

contained within the new system profile. This allows you to create multiple, nearly identical system profiles quickly and easily.

About this task

To copy a system profile using the HMC, follow these steps:

Procedure



1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **System Actions > Legacy > Manage System Profiles**.
4. Select the system profile and click **Actions > Copy**.
5. Enter the name that you want to use for the copy into **New profile name** and click **OK**.

Changing a system profile

You can change which partition profiles are included in a system profile using the Hardware Management Console (HMC).

Before you begin

Restriction: You cannot add logical partitions that use shared memory to system profiles.

About this task

To change a system profile using the HMC, follow these steps:

Procedure



1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **System Actions > Legacy > Manage System Profiles**.
4. Select the system profile that you want to change and click **Actions > Edit**.
5. In the **System Profile** window, select each partition profile that you want to remove from the system profile and click **Remove**.
6. For each partition profile that you want to add to the system profile, open the logical partition to which the partition profile belongs, select the partition profile, and click **Add**.
7. Click **OK**.

Validating a system profile

When you validate a system profile, the Hardware Management Console (HMC) compares the resources defined in the system profile with the resources available on the managed system. If the system profile requires more resources than are available on the managed system, a message is displayed on the HMC.

About this task

To validate a system profile using the HMC, follow these steps:

Procedure



1. In the navigation pane, click the **Resources** icon .

2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **System Actions > Legacy > Manage System Profiles**.
4. Select the system profile and click **Validate**.
5. When validation is complete, click **OK**.

Deleting a system profile

You can delete a system profile using the Hardware Management Console (HMC). This allows you to remove system profiles that you no longer use.

Before you begin

A system profile helps you activate or change the managed system from one complete set of logical partition configurations to another.

About this task

To delete a system profile using the HMC, follow these steps:

Procedure

1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **System Actions > Legacy > Manage System Profiles**.
4. Select the system profile and click **Actions > Delete**.
5. Click **Yes** to confirm.

Managing the resources of a shutdown logical partition

You can use the Hardware Management Console (HMC) command-line interface to manage the resources of a shutdown logical partition.

You can use the **chhwres** command to remove memory, processor, and I/O resources from a shutdown logical partition.

You can change other attributes of a shutdown logical partition by changing the logical partition profile and applying the changed profile to the logical partition. Complete the following steps from the HMC command line:

1. To change the profile of a shutdown logical partition, run the following command:

```
chsyscfg -r prof -m managed system -i attributes
```

Note: If the profile being changed is the last activated profile, you must use the *--force* option if synchronization of the current configuration to the profile is enabled for the logical partition.

2. To apply the changed profile to the shutdown logical partition, run the following command:

```
chsyscfg -r lpar -m managed system -o apply -n profile name
```

Managing logical partition resources dynamically

You can use the Hardware Management Console (HMC) to add, remove, or move processor, memory, and I/O resources between running logical partitions without restarting the logical partitions or the managed system.

Dynamic Platform Optimizer

POWER7 with firmware level FW760 or later, POWER8, or POWER9 processor-based servers supports the Dynamic Platform Optimizer (DPO) function. DPO is a hypervisor function that is initiated from the Hardware Management Console (HMC). DPO rearranges logical partition processors and memory on the system to improve the affinity between processors and memory of logical partitions. When DPO is running, mobility operations that target the system that is being optimized are blocked. Also, when DPO is running, many virtualization features are blocked. When a DPO operation is in progress and you want to dynamically add, remove, or move physical memory to or from running logical partitions, you must either wait for the DPO operation to complete or manually stop the DPO operation.

To help assess when DPO might be beneficial, you can use the HMC to determine affinity scores for the system and logical partitions by using the **lsmemopt** command. An affinity score is a measure of the processor-memory affinity on the system or for a partition. The score is a number in the range 0 - 100, 0 represents the worst affinity and 100 represents perfect affinity. Based on the system configuration, a score of 100 might not be attainable. A partition that has no processor and memory resources does not have an affinity score, and none is displayed for the score on the command line, when you run the **lsmemopt** command.

In addition to manually running DPO by using the **optmem** command, you can schedule DPO operations on POWER7 with firmware level FW760 or later, POWER8, or POWER9 processor-based servers. The HMC must be at Version 7.8.0 or later. The following conditions apply to the DPO operation:

- The current server affinity score of the managed system is less than or equal to the server affinity threshold that you provided.
- The affinity delta (which is the potential score minus the current score) of the managed system is greater than or equal to the affinity delta threshold of the server that you provided.

The scheduled operation sends a DPO report after the successful completion of a DPO operation, only if it is enabled in the **HMC Notifications**.

Querying affinity scores of a logical partition

On POWER7 or POWER9 processor-based servers with firmware at level FW780, or later, the HMC provides an additional flag with the **lsmemopt** command for querying the current affinity score and the potential affinity score of a logical partition.

About this task

Procedure

1. From the HMC command line, type the following command to query the current and potential logical partition affinity scores:

```
lsmemopt -m managed system -r lpar -o currscore | calcscore [-p partition-names | --id  
partition-IDs]  
[-x partition-names | --xid partition-IDs]
```

where:

- *currscore* queries the current affinity scores.
- *calcscore* queries the current and potential affinity scores.
- *-x partition-names* or *--xid partition-IDs* specifies the list of logical partitions or logical partition IDs that must not be affected by the optimization operation.
- *-p partition-names* or *--id partition-IDs* specifies the list of logical partitions or logical partition IDs that must be optimized.

The following example shows a sample output of the **lsmemopt** command when the **-o currscore** parameter is specified:

```
lpar_name=x,lpar_id=1,curr_lpar_score=25
```

The following example shows a sample output of the **lsmemopt** command when the **-o calcscore** parameter is specified:

```
lpar_name=x,lpar_id=1,curr_lpar_score=25,predicted_lpar_score=100
```

2. From the HMC command line, type the following command to query the system-wide affinity scores:

```
lsmemopt -m managed system -o currscore | calcscore [-p partition-names | --id partition-IDs]  
[-x partition-names | --xid partition-IDs]
```

where:

- **currscore** queries the current affinity scores.
- **calcscore** queries the current and potential affinity scores.
- **-x partition-names** or **--xid partition-IDs** specifies the list of logical partitions or logical partition IDs that must not be affected by the optimization operation.
- **-p partition-names** or **--id partition-IDs** specifies the list of logical partitions or logical partition IDs that must be optimized.

Scheduling Dynamic Platform Optimizer operations

Scheduled operation of the Dynamic Platform Optimizer (DPO) function is supported on POWER7 or POWER9 processor-based servers with firmware at level 7.6, or later. The Hardware Management Console (HMC) must be at Version 7.8.0 or later.

About this task

To schedule DPO operations by using the HMC, complete the following steps:

Procedure



1. In the navigation pane, click the **Resources** icon.
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **Actions > Schedule Operations**.
4. On the **Options** tab, click **New**.
5. Click **Monitor/Perform Dynamic Platform Optimize**.
6. Click **OK**.
7. On the Setup a Scheduled Operation page, click the **Date and Time** tab.
You can specify the date and time when the scheduled operation must start.
8. Click **Save**.
9. On the Setup a Scheduled Operation page, click the **Repeat** tab.
You can specify whether the scheduled operation is a single scheduled operation or a repeated scheduled operation. You can also specify the days of the week the operation must be performed, the interval and the number of repetitions. Click **Repeat Indefinitely** to perform the operation repeatedly for an indefinite period.
10. Click **Save**.
11. Setup a Scheduled Operation page, click the **Options** tab.
 - a) In the **Target of Operation** area, the system name and potential and current affinity scores are displayed.

The potential affinity score is a value in the range 0 -100 and it is queried from the HMC when the schedule operations option is selected. You can also use the **lsmemopt** command to get this value from the HMC command line. The current affinity score is a value in the range 0 -100 and it is queried from the HMC when the schedule operations option is selected. You can also use the **lsmemopt** command to get this value from the HMC command line.

- b) In the **Affinity Thresholds** area, you can specify a value in the range 0 -100 for the **Server Affinity Threshold** field.
- c) In the **Server Affinity Delta Threshold (Potential - Current)** field, enter a value.
- d) In the **Alert/Actions** area, when the email notification is not configured on the HMC, a message is displayed that informs you to configure the email notification. Click **Configure Management Console Notifications** to configure the email notifications.
- e) In the **Alert/Actions** area, when the email notification is configured on the HMC, click **Notify via an email of Server Affinity Alerts** to receive email notification alerts about DPO events.
- f) In the **Perform Dynamic Platform Optimization** area, click **Automatically Perform a Dynamic Platform Optimization (DPO)** to enable automatic DPO.



Attention: The DPO operation might automatically run constantly if the DPO does not cause the affinity to drop below either of the user-defined threshold values. This might impact system performance and block various virtualization functions. You can avoid setting the user-defined threshold values with a small interval when the auto-DPO option is enabled.

12. Click **Save**.

Starting and stopping a Dynamic Platform Optimizer operation

You can run the **optmem** command from the Hardware Management Console (HMC) command line on POWER7 or POWER9 processor-based servers with firmware at level FW760, or later, to start a Dynamic Platform Optimizer (DPO) operation or stop a DPO operation that is currently running.

Procedure

1. From the HMC command line, type the following command to start a DPO operation:

```
optmem -m managed-system -o start -t affinity [-p partition-names | --id partition-IDs]
[-x partition-names | --xid partition-IDs]
```

Where:

- *-x partition-names* or *--xid partition-IDs* specifies the list of logical partitions or logical partition IDs that must not be affected by the optimization operation.
- *-p partition-names* or *--id partition-IDs* specifies the list of logical partitions or logical partition IDs that must be optimized

2. To stop a currently running DPO operation, complete the following steps:

- a) From the HMC command line, type the following command to list the DPO operation that is currently running:

```
lsmemopt -m managed-system
```

- b) From the HMC command line, type the following command to stop the DPO operation:

```
optmem -m managed-system -o stop [--optid ID]
```

Where:

- *--optid* is an optional parameter that identifies the DPO operation to be canceled.
- *ID* is the value returned by the **lsmemopt** command.



Attention: Stopping a DPO operation before completion might worsen the affinity state of the system as compared to the affinity state of the system when the DPO operation was started.

Managing dedicated memory dynamically

You can add, remove, and move physical memory dynamically to and from running logical partitions that use dedicated memory by using the Hardware Management Console (HMC). This allows you to adjust the physical memory allocated to each logical partition that uses dedicated memory without having to shut down the logical partitions.

When a DPO operation is in progress and you want to dynamically add, remove, or move physical memory to, or from running logical partitions, you must either wait for the DPO operation to complete, or manually stop the DPO operation.

Dynamic memory changes on IBM i logical partitions affect the base memory pool of the logical partitions (*BASE pool). Private memory pools or shared memory pools are not affected. Dynamic memory changes cannot cause the amount of memory in the base pool to fall below the minimum amount of memory required in the base pool (as determined by the base storage minimum size (QBASPOOL) system value). If a dynamic memory change would cause the base pool to fall below this amount, the system releases excess memory pages only after keeping the minimum amount of memory required in the base pool.

To prevent any data loss during dynamic memory movement, the system first writes any data from memory pages to disk before making the memory pages available to another logical partition. Depending on the amount of memory you have requested to move, this might take some time.

Memory in each logical partition operates within its assigned minimum and maximum values. The full amount of memory that you assign to a logical partition might not be available for the logical partition to use. Static memory overhead that is required to support the assigned maximum memory affects the reserved or hidden memory amount. This static memory overhead also influences the minimum memory size of a logical partition.

Note:

- If resources are moved dynamically, the configuration change is temporary and is not reflected in the partition profile. This means that all configuration changes will be lost the next time the partition profile is activated. If you want to save your new logical partition configuration, either change the partition profile or save the logical partition configuration to a new partition profile.
- When dynamic logical partitioning tasks to add, remove, or move physical memory are run concurrently for a logical partition, the logical partition might not have the expected amount of physical memory after the concurrent tasks are complete. The logical partition might not have the expected amount of physical memory whether you specify the amount of physical memory that you want the logical partition to have after the dynamic logical partitioning task is complete, or you specify the amount of physical memory to be added to, removed from, or moved to or from the logical partition.

Adding dedicated memory dynamically

You can dynamically add physical memory to a running logical partition that uses dedicated memory using the Hardware Management Console (HMC). This allows you to increase the physical memory available to a logical partition that uses dedicated memory without having to shut down the logical partition.

Before you begin

A Linux logical partition supports the dynamic addition of memory resources only if the following conditions are met:

- A Linux distribution that supports the dynamic addition of memory resources is installed on the Linux logical partition. Distributions that support the dynamic addition of memory resources include SUSE Linux Enterprise Server 10, and later.
- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the [Service and productivity tools for Linux on POWER systems](#) website.

To add memory to a Linux logical partition that uses an earlier version of these distributions, you must shut down the Linux logical partition and reactivate the logical partition using a partition profile that specifies a greater amount of memory.

About this task

For more information about changing the memory settings, see [Changing memory settings](#).

Changing the Active Memory Expansion factor for AIX logical partitions

You can dynamically change the Active Memory Expansion factor for an AIX logical partition by using the Hardware Management Console (HMC). Changing the Active Memory Expansion factor for a logical partition increases or decreases the desired degree of expanded memory capacity for the logical partition.

Before you begin

You can change the Active Memory Expansion factor for logical partitions that use dedicated memory and logical partitions that use shared memory.

About this task

For more information about changing the memory settings, see [Changing memory settings](#).

Moving dedicated memory dynamically

You can dynamically move physical memory from one running logical partition that uses dedicated memory to another using the Hardware Management Console (HMC). This allows you to reassign physical memory directly to a logical partition that uses dedicated memory that needs additional physical memory.

Before you begin

You cannot dynamically move memory from a running Linux logical partition. To remove memory from a Linux logical partition, you must shut down the Linux logical partition and reactivate the logical partition using a partition profile that specifies a lesser amount of memory.

You can dynamically move memory to a running Linux only if the following conditions are met:

- A Linux distribution that supports the dynamic addition of memory resources is installed on the Linux logical partition. Distributions that support the dynamic movement of memory resources include Novell SUSE Linux Enterprise Server 10, and later.
- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the [Service and productivity tools for Linux on POWER systems](#) website.

To move memory to a Linux logical partition that uses an earlier version of these distributions, you must shut down the Linux logical partition and reactivate the logical partition using a partition profile that specifies a greater amount of memory.

About this task

For more information about changing the memory settings, see [Changing memory settings](#).

Removing dedicated memory dynamically

You can dynamically remove physical memory from a running AIX, IBM i, or Virtual I/O Server logical partition that uses dedicated memory using the Hardware Management Console (HMC). This allows you to reassign the physical memory to other logical partitions that use dedicated memory.

Before you begin

You cannot dynamically remove memory from a running Linux logical partition. To remove memory from a Linux logical partition, you must shut down the logical partition and reactivate the logical partition using a partition profile that specifies a lesser amount of memory.

About this task

For more information about changing the memory settings, see [Changing memory settings](#).

Managing shared memory dynamically

You can dynamically add and remove logical memory and I/O entitled memory to and from a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) using the Hardware Management Console (HMC).

About this task

Dynamic memory changes on IBM i logical partitions affect the base memory pool of the logical partitions (*BASE pool). Private memory pools or shared memory pools are not affected. Dynamic memory changes cannot cause the amount of memory in the base pool to fall below the minimum amount of memory required in the base pool (as determined by the base storage minimum size (QBASPOOL) system value). If a dynamic memory change would cause the base pool to fall below this amount, the system releases excess memory pages only after keeping the minimum amount of memory required in the base pool.

Note:

- If resources are moved dynamically, the configuration change is temporary and is not reflected in the partition profile. This means that all configuration changes will be lost the next time the partition profile is activated. If you want to save your new logical partition configuration, either change the partition profile or save the logical partition configuration to a new partition profile.
- When dynamic logical partitioning tasks to add or remove shared memory are run concurrently for a logical partition, the logical partition might not have the expected amount of shared memory after the concurrent tasks are complete. The logical partition might not have the expected amount of shared memory whether you specify the amount of shared memory that you want the logical partition to have after the dynamic logical partitioning task is complete, or you specify the amount of shared memory to be added to or removed from the logical partition.

Adding and removing logical memory dynamically to and from a shared memory partition

You can dynamically add and remove logical memory to and from a running logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) using the Hardware Management Console (HMC). This allows you to increase and decrease the logical memory assigned to the shared memory partition without having to shut down the logical partition.

Before you begin

A Linux shared memory partition supports the dynamic addition and removal of logical memory resources only if the DynamicRM tool package is installed on the Linux shared memory partition. To download the DynamicRM tool package, see the [Service and productivity tools for Linux on POWER systems](#) website.

To dynamically add and remove logical memory to and from a running logical partition using the HMC, you must be a super administrator, service representative, product engineer, or operator of the HMC.

About this task

For more information about changing the memory settings, see [Changing memory settings](#).

Adding and removing I/O entitled memory dynamically to and from a shared memory partition

You can dynamically add and remove I/O entitled memory to and from a running logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) using the Hardware Management Console (HMC). This allows you to increase and decrease the maximum amount of physical memory that is assigned to the shared memory partition for its I/O devices without having to shut down the shared memory partition.

Before you begin

A Linux shared memory partition supports the dynamic addition and removal of I/O entitled memory resources only if the DynamicRM tool package is installed on the Linux shared memory partition. To download the DynamicRM tool package, see the [Service and productivity tools for Linux on POWER systems](#) website.

You can increase the amount of I/O entitled memory that is assigned to a shared memory partition when the sum of I/O entitled memory that is assigned to all shared memory partitions in the shared memory pool is less than the size of the shared memory pool minus the required amount of reserved firmware memory. If there is not enough physical memory in the shared memory pool by which to increase the I/O entitled memory to the amount specified, you can release to the hypervisor the physical memory that is currently assigned to other shared memory partitions that are shut down. The hypervisor can then assign the released physical memory to the shared memory partition that needs more I/O entitled memory.

You can decrease the amount of I/O entitled memory that is assigned to a shared memory partition only when the shared memory partition requires less physical memory for its I/O devices than the amount of I/O entitled memory that is assigned to the shared memory partition. For example, you assign 128 MB of I/O entitled memory to a shared memory partition. The shared memory partition requires a minimum of 64 MB for its I/O devices. Thus, you can decrease the I/O entitled memory that is assigned to the shared memory partition by up to 64 MB. For instructions about how to view the assigned, minimum, optimal, and maximum I/O entitled memory used by a shared memory partition, see [“Determining the I/O entitled memory for a shared memory partition”](#) on page 219.

To dynamically add and remove I/O entitled memory to and from a running shared memory partition using the HMC, you must be a super administrator, service representative, product engineer, or operator of the HMC.

About this task

For more information about changing the memory settings, see [Changing memory settings](#).

Results

If you want to later change the I/O entitled memory mode back to the auto mode so that the HMC automatically adjusts the I/O entitled memory for the shared memory partition when you add or remove virtual adapters, repeat this procedure and select **Auto**. Alternatively, you can restart the shared memory partition. When you restart a shared memory partition, the I/O entitled memory mode is set to the auto mode regardless of what the I/O entitled memory mode was set to before you restarted the shared memory partition.

Changing the Active Memory Expansion factor for AIX logical partitions

You can dynamically change the Active Memory Expansion factor for an AIX logical partition by using the Hardware Management Console (HMC). Changing the Active Memory Expansion factor for a logical partition increases or decreases the desired degree of expanded memory capacity for the logical partition.

Before you begin

You can change the Active Memory Expansion factor for logical partitions that use dedicated memory and logical partitions that use shared memory.

About this task

For more information about changing the memory settings, see [Changing memory settings](#).

Managing processor resources dynamically

You can dynamically add, remove, and move processor resources to and from running logical partitions using the Hardware Management Console (HMC). This allows you to adjust the processor resources allocated to each logical partition without having to shut down the logical partitions.

The ability to move processor resources dynamically becomes important when you need to adjust to changing workloads. Processor resources can be moved based on the minimum and maximum values that you created for the partition profile. You can move processor resources as long as the processor resources for each logical partition remains within the range specified by the minimum and maximum values for the logical partition. If the managed system uses more than one shared processor pool, you must also ensure that the number of processors used in each shared processor pool is less than or equal to the maximum number of processing units specified for each shared processor pool.

Note:

- If resources are moved dynamically, the configuration change is temporary and is not reflected in the partition profile. This means that all configuration changes will be lost the next time the partition profile is activated. If you want to save your new logical partition configuration, either change the partition profile or save the logical partition configuration to a new partition profile.
- When dynamic logical partitioning tasks to add, remove, or move processor resources are run concurrently for a logical partition, the logical partition might not have the expected number of processor resources after the concurrent tasks are complete. The logical partition might not have the expected number of processor resources whether you specify the number of processor resources that you want the logical partition to have after the dynamic logical partitioning task is complete, or you specify the number of processor resources to be added to, removed from, or moved to or from the logical partition.

Adding processor resources dynamically

You can dynamically add processor resources to a running logical partition using the Hardware Management Console (HMC). This allows you to increase the processing capacity of a running logical partition without having to shut down the logical partition.

Before you begin

A Linux logical partition supports the dynamic addition of processor resources only if the following conditions are met:

- A Linux distribution that supports dynamic partitioning is installed on the Linux logical partition. Distributions that support dynamic partitioning include SUSE Linux Enterprise Server 9 and later versions.
- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the [Service and productivity tools for Linux on POWER systems](#) website.

About this task

For more information about changing the processor settings, see [Changing processor settings](#).

Moving processor resources dynamically

You can dynamically move processor resources from one running logical partition to another using the Hardware Management Console (HMC). This allows you to reassign processor resources directly to a logical partition that needs additional processor resources.

Before you begin

A Linux logical partition supports the dynamic movement of processor resources only if the following conditions are met:

- A Linux distribution that supports dynamic partitioning is installed on the Linux logical partition. Distributions that support dynamic partitioning include SUSE Linux Enterprise Server 9 and later versions.

- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the [Service and productivity tools for Linux on POWER systems](#) website.

About this task

For more information about changing the processor settings, see [Changing processor settings](#).

Removing processor resources dynamically

You can dynamically remove processor resources from a running logical partition using the Hardware Management Console (HMC). This allows you to reassign the processor resources to other logical partitions.

Before you begin

A Linux logical partition supports the dynamic removal of processor resources only if the following conditions are met:

- A Linux distribution that supports dynamic partitioning is installed on the Linux logical partition. Distributions that support dynamic partitioning include SUSE Linux Enterprise Server 9 and later versions.
- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the [Service and productivity tools for Linux on POWER systems](#) website.

About this task

For more information about changing the processor settings, see [Changing processor settings](#).

Managing physical I/O devices and slots dynamically

You can dynamically add, remove, and move physical I/O devices and slots to and from running logical partitions using the Hardware Management Console (HMC). This allows logical partitions to share infrequently used I/O devices (such as optical disk drives).

Logical partitions can have desired or required I/O devices or slots. When you specify that an I/O device or slot is desired, this means either that the I/O device or slot is meant to be shared with other logical partitions, or that the I/O device or slot is optional. When you specify that an I/O device or slot is required (or dedicated), then you cannot activate the logical partition if the I/O device or slot is unavailable or in use by another logical partition.

Note: If resources are moved dynamically, the configuration change is temporary and is not reflected in the partition profile. This means that all configuration changes will be lost the next time the partition profile is activated. If you want to save your new logical partition configuration, either change the partition profile or save the logical partition configuration to a new partition profile.

Adding physical I/O devices and slots dynamically

You can dynamically add a physical I/O slot (and the adapter and devices that are connected to that slot) to a running logical partition using the Hardware Management Console (HMC). This allows you to add I/O capabilities to a running logical partition without having to shut down the logical partition.

Before you begin

A Linux logical partition supports the dynamic addition of physical I/O slots only if the following conditions are met:

- A Linux distribution that supports dynamic partitioning is installed on the Linux logical partition. Distributions that support dynamic partitioning include SUSE Linux Enterprise Server 9, and later.
- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the [Service and productivity tools for Linux on POWER systems](#) website.

You cannot add physical I/O devices and slots to logical partitions that use shared memory. You can assign only virtual adapters to logical partitions that use shared memory.

About this task

For more information about managing physical I/O adapters, see [Managing physical I/O adapters](#).

Moving physical I/O devices and slots dynamically

You can dynamically move a physical I/O slot (and the adapter and devices that are connected to that slot) from one running logical partition to another using the Hardware Management Console (HMC). This allows you to share a physical I/O device, such as a DVD drive, among many logical partitions.

Before you begin

Before you begin, vary off any devices that are attached to the managed system through the physical I/O slot that you want to move. You can vary off devices by using operating system commands.



Attention: The dynamic movement of a physical I/O slot that controls disk drives can cause unpredictable results, such as logical partition failure or loss of data.

A Linux logical partition supports the dynamic movement of physical I/O slots only if the following conditions are met:

- A Linux distribution that supports dynamic partitioning is installed on the Linux logical partition. Distributions that support dynamic partitioning include SUSE Linux Enterprise Server 9, and later.
- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the [Service and productivity tools for Linux on POWER systems](#) website.

You cannot dynamically move physical I/O devices and slots to logical partitions that use shared memory. You can assign only virtual adapters to logical partitions that use shared memory.

About this task

For more information about managing physical I/O adapters, see [Managing physical I/O adapters](#).

Removing physical I/O devices and slots dynamically

You can dynamically remove a physical I/O slot and the adapter and devices that are connected to that slot from a running logical partition using the Hardware Management Console (HMC). This allows you to reassign the physical I/O slot to other logical partitions.

Before you begin

Before you begin, vary off any devices that are attached to the managed system through the physical I/O slot that you want to remove. You can vary off devices using operating system commands.



Attention: The dynamic removal of a physical I/O slot that controls disk drives can cause unpredictable results, such as logical partition failure or loss of data.

A Linux logical partition supports the dynamic removal of physical I/O slots only if the following conditions are met:

- A Linux distribution that supports dynamic partitioning is installed on the Linux logical partition. Distributions that support dynamic partitioning include SUSE Linux Enterprise Server 9, and later.
- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the [Service and productivity tools for Linux on POWER systems](#) website.

About this task

For more information about managing physical I/O adapters, see [Managing physical I/O adapters](#).

Managing virtual adapters dynamically

You can dynamically add and remove virtual adapters to and from running logical partitions using the Hardware Management Console (HMC).

About this task

Tasks that are related to managing the virtual adapters such as adding a virtual adapter or removing an adapter are a part of managing the logical partition. When you perform tasks on virtual storage (vSCSI, virtual Fibre Channel, virtual optical device, virtual network, virtual NICs), the virtual adapters operations are handled automatically. For more information about virtual storage, see [Managing virtual storage](#). For more information about virtual networks, see [Managing virtual networks](#). For more information about virtual Network Interface Controllers (vNICs), see [Managing virtual Network Interface Controllers](#).

Note: If resources are moved dynamically, the configuration change is temporary and is not reflected in the partition profile. This means that all configuration changes will be lost the next time the partition profile is activated. If you want to save your new logical partition configuration, either change the partition profile or save the logical partition configuration to a new partition profile.

Enabling and disabling SR-IOV logical ports

You can enable or disable single root I/O virtualization (SR-IOV) logical ports by using the Hardware Management Console (HMC). The HMC must be at Version 9.1.0, or later.

Before you begin

You must be a super administrator to complete this task.

About this task

- To disable an SR-IOV logical port, type the following command from the HMC command line:

```
chhwres -m <managed-system> -r sriov --rsubtype logport -o d [-p <partition-name> | --id <partition-ID>]
-a "adapter_id=<adapter-id>,logical_port_id=<logical-port-id>"
```

The logical partition can be either in the running or shut down state.

- To enable an SR-IOV logical port, type the following command from the HMC command line:

```
chhwres -m <managed-system> -r sriov --rsubtype logport -o e [-p <partition-name> | --id <partition-ID>]
-a "adapter_id=<adapter-id>,logical_port_id=<logical-port-id>"
```

The logical partition can be either in the running or shut down state.

- To verify whether SR-IOV logical ports are enabled or disabled, type the following command from the HMC command line:

```
lshwres -m <managed-system> -r sriov --rsubtype logport [--filter "<filter-data>"] -level <type>
[-F [<attribute-names>] [--header]]
```

If the value of the *is_disabled* attribute is 0, the SR-IOV logical ports are enabled. If the value of the *is_disabled* attribute is 1, the SR-IOV logical ports are disabled.

Related information

[chhwres command](#)

[lshwres command](#)

Managing SR-IOV logical ports dynamically

You can dynamically add, edit, and remove single root I/O virtualization (SR-IOV) logical ports to and from running logical partitions by using the Hardware Management Console (HMC).

Adding a single root I/O virtualization logical port to a logical partition dynamically

You can dynamically add a single root I/O virtualization (SR-IOV) logical port to a running logical partition by using the Hardware Management Console (HMC).

About this task

For more information about SR-IOV logical port settings, see [SR-IOV logical port settings](#).

When the HMC is at Version 9.1.0, or later, you can use the *max_capacity* attribute of the **chhwres** command to specify the maximum capacity value for the SR-IOV logical port when you are adding an SR-IOV logical port to a logical partition.

When the HMC is at Version 9.1.940, or later, and when the firmware is at level FW940, or later, you can use the *migratable* attribute of the **chhwres** command to add a migratable SR-IOV logical port to a logical partition.

Note: When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

Related information

[chhwres command](#)

Viewing migratable SR-IOV logical ports and SR-IOV backup virtual devices

You can view the list of single root I/O virtualization (SR-IOV) logical ports and SR-IOV backup virtual devices that can be migrated by using the Hardware Management Console (HMC) command-line interface.

About this task

When the HMC is at Version 9.1.940, or later, and when the firmware is at level FW940, or later, you can run the **lshwres** command from the HMC command line to view the migratable attribute of the single root I/O virtualization (SR-IOV) logical port and SR-IOV backup virtual devices of migratable SR-IOV logical ports.

Note: When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

For more information about SR-IOV logical port settings, see [SR-IOV logical port settings](#).

Related information

[lshwres command](#)

Modifying a single root I/O virtualization logical port that is assigned to a logical partition dynamically

You can modify a single root I/O virtualization (SR-IOV) logical port that is assigned to a running logical partition by using the Hardware Management Console (HMC).

About this task

When the HMC is at Version 9.1.940, or later, and when the firmware is at level FW940, or later, you can use the **chhwres** command to change the Port VLAN ID (PVID), allowed VLANs, and allowed operating system MAC addresses of a configured and migratable SR-IOV logical port. Additionally, the HMC checks whether the backup device of the SR-IOV logical port is a virtual Network Interface Controller (vNIC). If the backup device is a vNIC, the changes that were applied to the SR-IOV logical port are also applied to the vNIC and to the backing device of the vNIC.

Note: When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

For more information about SR-IOV logical port settings, see [SR-IOV logical port settings](#).

Related information

[chhwres command](#)

Removing a single root I/O virtualization logical port from a logical partition dynamically

You can dynamically remove a single root I/O virtualization (SR-IOV) logical port from a running logical partition by using the Hardware Management Console (HMC).

About this task

For more information about SR-IOV logical port settings, see [SR-IOV logical port settings](#).

Note:

- When the HMC is at Version 9.1.940, or later, and when the firmware is at level FW940, or later, you cannot remove a backup virtual device when an SR-IOV logical port is associated with it.
- When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

Creating a profile with migratable single root I/O virtualization logical ports

You can create a profile with migratable single root I/O virtualization (SR-IOV) logical ports by using the Hardware Management Console (HMC).

About this task

When the HMC is at Version 9.1.940, or later, and when the firmware is at level FW940, or later, you can create a profile with migratable SR-IOV logical ports by using the *migratable* attribute of the **mksyscfg** command. You must also specify the *backup_veth_vnetwork* attribute for the virtual Ethernet backup devices.

Note: When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

For more information about SR-IOV logical port settings, see [SR-IOV logical port settings](#).

Related information

[mksyscfg command](#)

Recovering a migratable single root I/O virtualization logical port

When the Hardware Management Console (HMC) is at Version 9.1.940, or later, and when the firmware is at level FW940, or later, if a virtual I/O device is specified as the backup device for a migratable single root I/O virtualization (SR-IOV) logical port, the migratable SR-IOV logical port might not be available after running the migration operation was run by using the *--migsriov 2* option of the **migrldpar** command. You can use the *recover* option of the **chhwres** command to recover the SR-IOV logical port. When the HMC is at Version 9.1.940.x, and when the firmware is at level FW940, the Migratable option for the Hybrid Network Virtualization capability is available as a Technology Preview only and is not intended for production deployments. However, when the HMC is at Version 9.1.941.0, or later, and when the firmware

is at level FW940.10, or later, the Migratable option for the Hybrid Network Virtualization capability is supported.

About this task

For more information about SR-IOV logical port settings, see [SR-IOV logical port settings](#).

Related information

[migrpar command](#)

Managing 5250 CPW dynamically

You can dynamically add, remove, and move 5250 commercial processing workload (5250 CPW) to and from running logical partitions using the Hardware Management Console (HMC).

5250 CPW is the capacity to perform 5250 online transaction processing (5250 OLTP) tasks on IBM i logical partitions. On certain servers, you can assign a percentage of the total 5250 CPW available on the managed system to each IBM i logical partition. The ability to assign 5250 CPW to IBM i logical partitions is available only for Express Configurations and Value Editions.

5250 CPW can be moved based on the desired, minimum, and maximum percentages you created for the partition profile. The desired 5250 CPW percentage you establish is the amount of 5250 CPW that you get if you do not overcommit the available 5250 CPW. The minimum and maximum values enable you to establish a range within which you can dynamically move the 5250 CPW.



Attention: If resources are moved dynamically, the configuration change is temporary and is not reflected in the partition profile. This means that all configuration changes will be lost the next time the partition profile is activated. If you want to save your new logical partition configuration, you should change the partition profile or save the logical partition configuration to a new partition profile.

Adding 5250 CPW for IBM i logical partitions dynamically

You can dynamically add 5250 commercial processing workload (5250 CPW) to a running IBM i logical partition using the Hardware Management Console (HMC). This allows you to increase the ability of the IBM i logical partition to run 5250 online transaction processing (5250 OLTP) tasks.

Before you begin

This procedure applies only to Express Configurations and Value Editions, which provide a fixed amount of processing capability for 5250 OLTP tasks.

About this task

To add 5250 CPW to a running IBM i logical partition using the HMC, follow these steps:

Procedure

1. In the navigation pane, open **Systems Management > Servers**, and click the managed system on which the logical partition resides.
2. In the work pane, select the logical partition, click the **Tasks** button, and click **Dynamic Partitioning > Processor > Add or Remove**.
3. Enter the amounts of 5250 CPW that you want the logical partition to have into the **5250 CPW (percent)** field in the **Current** column.
4. Adjust the settings in the **Options** area if necessary.
You might need to increase the value in the **Timeout (minutes)** field to allow enough time for the HMC to complete the operation. (These settings relate to how the managed system adds 5250 CPW dynamically. These settings are not retained after the addition is completed.)
5. Click **OK**.

Moving 5250 CPW for IBM i logical partitions dynamically

You can dynamically move 5250 commercial processing workload (5250 CPW) from one running IBM i logical partition to another using the Hardware Management Console (HMC). This allows you to use the limited amount of 5250 CPW that is available on your managed system efficiently.

Before you begin

This procedure applies only to Express Configurations and Value Editions, which provide a fixed amount of processing capability for 5250 online transaction processing (5250 OLTP) tasks.

About this task

To move 5250 CPW from one running IBM i logical partition to another using the HMC, follow these steps:

Procedure

1. In the navigation pane, open **Systems Management > Servers**, and click the managed system on which the logical partitions reside.
2. In the work pane, select the logical partition from which you want to move 5250 CPW, click the **Tasks** button, and click **Dynamic Partitioning > Processor > Move**.
3. Enter the amounts of 5250 CPW that you want to move into the **5250 CPW (percent)** field in the **To move** column.
4. Select the logical partition to which you want to move 5250 CPW in **Select Destination Partition**.
5. Adjust the settings in the **Options** area if necessary.
You might need to increase the value in the **Timeout (minutes)** field to allow enough time for the HMC to complete the operation. (These settings relate to how the managed system moves 5250 CPW dynamically. These settings are not retained after the move is completed.)
6. Click **OK**.

Removing 5250 CPW for IBM i logical partitions dynamically

You can dynamically remove 5250 commercial processing workload (5250 CPW) dynamically from a running IBM i logical partition using the Hardware Management Console (HMC). This allows you to make 5250 CPW available for assignment to other IBM i logical partitions on the managed system.

Before you begin

This procedure applies only to Express Configurations and Value Editions, which provide a fixed amount of processing capability for 5250 online transaction processing (5250 OLTP) tasks.

About this task

To remove 5250 CPW from a running IBM i logical partition using the HMC, follow these steps:

Procedure

1. In the navigation pane, open **Systems Management > Servers**, and click the managed system on which the logical partition resides.
2. In the work pane, select the logical partition, click the **Tasks** button, and click **Dynamic Partitioning > Processor > Add or Remove**.
3. Enter the amounts of 5250 CPW that you want the logical partition to have into the **5250 CPW (percent)** field in the **Current** column.
4. Adjust the settings in the **Options** area if necessary.
You might need to increase the value in the **Timeout (minutes)** field to allow enough time for the HMC to complete the operation. (These settings relate to how the managed system removes 5250 CPW dynamically. These settings are not retained after the removal is completed.)
5. Click **OK**.

Scheduling the movement of resources to and from logical partitions

You can use the Hardware Management Console (HMC) to schedule the movement of dedicated memory, logical memory, dedicated processors, shared processors, and I/O devices between running logical partitions on a managed system. This allows you to move resources between running logical partitions without user intervention.

About this task

To schedule the movement of resources to or from a running logical partition using the HMC, follow these steps:

Procedure



1. In the navigation pane, click the **Resources** icon.
 2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
 3. In the work pane, select the logical partition and click **Actions > View Partition Properties**.
 4. Click **Partition Actions > Schedule Operations**.
 5. Click **Options** and click **New**.
 6. Select **Dynamic Reconfiguration**, and click **OK**.
 7. Select the date and time on which you want the movement to occur.
 8. Select the **Options** tab and select the resource type (I/O, memory, or processor), the type of movement (**Add**, **Remove**, or **Move to**), the destination logical partition (if you are moving resources to another logical partition), and the quantity (in processors or in megabytes) or the I/O slot that you want to move.
- Note:** You can add or remove logical memory to or from a logical partition. You cannot move logical memory from one logical partition to another logical partition.
9. If you want the operation to be repeated, select the **Repeat** tab and specify how you want the operation to be repeated.
 10. Click **Save**.
 11. When the message dialog displays, click **OK** to continue.

Results

When this procedure is completed, the managed system is set to perform the dynamic partitioning task at the date and time that you specify.

Saving the logical partition configuration to a partition profile

You can save the current configuration of a logical partition to a new partition profile using the Hardware Management Console (HMC). Use this procedure if you change the configuration of a logical partition using dynamic partitioning and you do not want to lose the changes when you reactivate the logical partition. This procedure allows you to save the changed configuration to a new partition profile instead of having to enter the changed resource allocations manually.

Before you begin

You can perform this procedure at any time after you initially activate a logical partition.

About this task

You can perform this procedure on active logical partitions and on logical partitions that are shut down. In either of these cases, the HMC reads the logical configuration that is stored for the logical partition in the server firmware and saves this logical configuration to the specified partition profile. For active logical partitions, the logical configuration that is stored in the server firmware is the current logical configuration

of the logical partition. For logical partitions that are shut down, the logical configuration that is stored in the server firmware is the logical configuration at the time that you shut down the logical partition. Regardless of the state of the logical partition at the time that you perform this procedure, the procedure allows you to save the dynamic partitioning changes to a partition profile and use the partition profile to reactivate the logical partition without losing those changes.

After you shut down a logical partition, other logical partitions can use the resources that were used by that logical partition when the logical partition was active. Therefore, the resources available on the managed system might not support the logical partition configuration that is stored in the server firmware for the inactive logical partition. After you save the logical configuration of a logical partition that is shut down, verify that the resources available on the managed system can support the logical partition configuration that you saved to a partition profile.

When you save the logical configuration to a new partition profile, the desired amounts of memory, processors, processing units, and virtual processors in the new partition profile are set to the current amounts from the logical configuration. The minimum and maximum amounts of memory, processors, processing units, and virtual processors in the new partition profile are set to the minimum and maximum amounts from the logical configuration. For example, you start a logical partition using a partition profile that specifies a minimum of 512 MB of dedicated memory, a maximum of 2 GB of dedicated memory, and 1 GB as the desired amount of dedicated memory. The managed system has over 1 GB of physical memory available, so the logical partition has 1 GB of physical memory when it starts. You then add 1 GB of physical memory to the logical partition for a total of 2 GB of physical memory. If you shut down the logical partition and then save the logical configuration, the resulting partition profile specifies a minimum of 512 MB of dedicated memory, a maximum of 2 GB of dedicated memory, and 2 GB as the desired amount of dedicated memory.

The physical and virtual I/O devices that are set as required in the active partition profile are saved as required devices in the new partition profile. The physical and virtual I/O devices that are set as desired in the active partition profile or that were added to the logical partition through dynamic partitioning are saved as desired devices in the new partition profile. The partition workload group on the logical partition (if any) is saved as the partition workload group on the new partition profile.

To save the current configuration of a logical partition to a new partition profile using the HMC, complete the following:

Procedure



1. In the navigation pane, click the **Resources** icon .
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions > Profiles > Save Current Configuration**.
4. Enter the name of the new partition profile into **New profile** and click **OK**.

What to do next

After you save the logical configuration to a new partition profile, verify that the new partition profile is set correctly. In particular, verify that the required and desired settings are set correctly for your I/O devices. By default, physical and virtual I/O devices that are added to the logical partition using dynamic partitioning are saved as desired devices in the new partition profile. If you want any of these I/O devices to be required, you must change the partition profile so that the I/O device is required.

Managing virtual resources for Virtual I/O Server logical partitions by using the HMC

Use the Hardware Management Console (HMC) to manage virtual storage that is associated with Virtual I/O Server logical partitions.

Changing a virtual disk for a VIOS logical partition by using the HMC

You can use the Hardware Management Console (HMC) to view the properties of the virtual disks on your managed system, as well as to start virtual disk management tasks.

About this task

Virtual disks are also known as logical volumes. To assign the virtual disk to a client partition, ensure that the client partition owns one or more virtual SCSI adapters and that the Virtual I/O Server (VIOS) owns corresponding virtual SCSI adapters that host the client adapter.

To change a virtual disk, be sure you meet the following requirements:

- The HMC must be at version 7.7.4, or later.
- The VIOS must be at version 2.2.1.0, or later.
- Ensure that there is a resource monitoring and control connection between the HMC and the VIOS.

To view and to change virtual disks, complete the following steps in the HMC:

Procedure



1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the **PowerVM** area, click **Virtual Storage**. The Virtual Storage page opens with the VIOS partitions listed in a table, in the Virtual Storage Management tab.
4. Select a VIOS and click **Action > Manage Virtual Storage**.
5. Click the **Virtual Disks** tab to display a list of virtual disks on the managed system.
6. Select the virtual disk from the table that you want to change.
If a virtual disk is defined as a paging space device and is assigned to a shared memory pool, it is dedicated to providing this function and is no longer available for any other purpose. Consequently, such a virtual disk is not listed here.
7. From the **Select Action** menu bar of the Virtual Disks table, select the storage management task you want to perform:
 - **Properties** to view the properties of the selected virtual disks.
 - **Extend** to add storage capacity to the selected virtual disks.
 - **Delete** to delete the selected virtual disk and make the storage resources that belonged to that virtual disk available to other virtual disks.
 - **Modify assignment** to change the logical partition to which the selected virtual disk is assigned, or to set the selected virtual disk so it is not assigned to any logical partitions.

Changing an optical device for a VIOS logical partition by using the Hardware Management Console

You can use the Hardware Management Console to view and to change physical optical devices and virtual optical media.

About this task

You can add optical devices to, or remove optical devices from, any logical partition, whether or not the logical partition is active. If you remove an optical device from an active logical partition, the Hardware Management Console prompts you to confirm the removal before removing the optical device. To assign an optical device to a client partition, ensure that the client partition owns one or more virtual SCSI adapters and that the VIOS owns corresponding virtual SCSI adapters that host the client adapter.

To change virtual optical media, be sure you meet the following requirements:

- The Hardware Management Console must be at version 7 release 3.4.2 or later.
- The Virtual I/O Server must be at version 2.1.1.0 or later.
- Ensure that there is a resource monitoring and control connection between the Hardware Management Console and the Virtual I/O Server.
- Verify that a virtual media library exists before you manage, create, or assign virtual optical devices.

To view and change optical devices, complete the following steps in the Hardware Management Console:

Procedure



1. In the navigation pane, click the **Resources** icon.
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the **PowerVM** area, click **Virtual Storage**. The Virtual Storage page opens with the VIOS partitions listed in a table, in the Virtual Storage Management tab.
4. Select a VIOS and click **Action > Manage Virtual Storage**.
5. Select a Virtual I/O Server logical partition.
6. Click the **Optical Devices** tab.
7. To change the logical partition assignment for a physical optical device, complete the following steps.
(You cannot assign a physical optical device to an IBM i logical partition. An IBM i logical partition must use virtual optical devices instead.)
 - a) From the Physical Optical Devices table, select the optical device that you want to change and click **Modify assignment**.
The Modify Physical Optical Device Assignment page is displayed.
 - b) Either change the logical partition to which the optical device is assigned, or set the optical device so it is not assigned to any logical partition, and click **OK**.
The list of optical devices reflects the changes you made.
8. To change virtual optical media, click one of the following tasks in the Virtual Optical Media section:
 - **Create/Extend Library** to extend the size of the media library.
 - **Delete Library** to delete the media library and the files within the library.
 - **Add Media** to add an optical media file to the media library and make it available for assignment to a partition.
 - **Modify partition assignment** to change the partition assignment for a media file by changing the virtual optical device to which a media file is assigned. You can assign read-only media to more than one partition.
 - **Delete** to delete the selected media files from the media library.

Changing a storage pool for a VIOS logical partition by using the HMC

You can use the Hardware Management Console (HMC) to extend a storage pool, to reduce or remove a storage pool, and to assign a storage pool as the default storage pool for the managed system.

About this task

To view and change storage pools, be sure you meet the following requirements:

- The Hardware Management Console must be at version 7 release 3.4.2 or later.
- The Virtual I/O Server must be at version 2.1.1.0 or later.
- Ensure that there is a resource monitoring and control connection between the Hardware Management Console and the Virtual I/O Server.

To view and to change storage pools, complete the following steps in the Hardware Management Console:

Procedure



1. In the navigation pane, click the **Resources** icon.
 2. Click **All Systems**. The **All Systems** page is displayed.
 3. In the **PowerVM** area, click **Virtual Storage**. The Virtual Storage page opens with the VIOS partitions listed in a table, in the Virtual Storage Management tab.
 4. Select a VIOS and click **Action > Manage Virtual Storage**.
 5. Select a Virtual I/O Server logical partition.
 6. Click the **Storage Pools** tab to display a list of storage pools defined for the managed system.
 7. Select the storage pool from the table that you want to change.
 8. From the **Select Action** menu bar of the Storage Pools table, select the storage management task that you want to perform:
 - **Properties** to view the properties of the selected storage pool.
 - **Extend** to add storage capacity to the selected storage pool. To extend logical volume-based storage pools, add physical volumes to the storage pool. To extend file-based storage pools, add space from the parent storage pool to the file-based storage pool.
- Note:** You cannot add a physical volume to a storage pool if it is already assigned to a partition.
- **Reduce** to reduce the size of the selected storage pool. To reduce logical volume-based storage pools, remove physical volumes from the storage pool. To reduce the file-based storage pool, the storage pool is deleted.



Attention: Reducing a storage pool that contains virtual disks could potentially destroy data stored on the virtual disks.

Changing a physical volume for a VIOS logical partition by using the HMC

You can use the Hardware Management Console (HMC) to view the properties of the physical volumes on your managed system, as well as to start physical volume management tasks.

About this task

A physical volume can be a hard disk or a logical device on a storage area network (SAN). You can either assign a physical volume directly to a logical partition, or you can add a physical volume to a storage pool and create virtual disks to assign to logical partitions from the storage pool.

To change physical volumes, be sure you meet the following requirements:

- The Hardware Management Console must be at version 7 release 3.4.2 or later.
- The Virtual I/O Server must be at version 2.1.1.0 or later.

- Ensure that there is a resource monitoring and control connection between the Hardware Management Console and the Virtual I/O Server.

To view and to modify physical volumes, complete the following steps in the Hardware Management Console:

Procedure



1. In the navigation pane, click the **Resources** icon.
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the **PowerVM** area, click **Virtual Storage**. The Virtual Storage page opens with the VIOS partitions listed in a table, in the Virtual Storage Management tab.
4. Select a VIOS and click **Action > Manage Virtual Storage**.
5. Select a Virtual I/O Server logical partition.
6. Click the **Physical Volumes** tab to display a list of physical volumes on the managed system.
7. Select the physical volume from the table that you want to change.
If a physical volume is defined as a paging space device and is assigned to a shared memory pool, it is dedicated to providing this function and is not available for any other purpose. Consequently, such a physical volume is not listed here.
8. From the **Select Action** menu bar of the Physical Volumes table, select the storage management task that you want to perform:
 - **Properties** to view or change the properties of the selected physical volume.
 - **Modify partition assignment** to change the logical partition to which the selected physical volume is assigned, or to set the physical volume so it is not assigned to any logical partition.
 - **Add to storage pool** to add the selected physical volume to a storage pool.
 - **Remove from storage pool** to remove the selected physical volume from the selected storage pool.

Changing virtual Fibre Channel for a Virtual I/O Server by using the HMC

You can use the Hardware Management Console (HMC) to dynamically manage virtual Fibre Channel on your managed system and the partition connections for the associated physical Fibre Channel ports. Assigning one or more physical ports to a logical partition enables the partition to communicate with storage devices in a storage area network (SAN). Configuring this type of storage resource is available only when the system supports the use of virtual Fibre Channel adapters and has a physical Fibre Channel adapter installed and connected that supports N_Port ID Virtualization (NPIV) ports.

Before you begin

To assign the virtual Fibre Channel adapter to a physical port, ensure that the client logical partition owns one or more virtual Fibre Channel adapters and that the Virtual I/O Server owns corresponding virtual Fibre Channel adapters to host the client adapter.

To change a port connection assignment for a logical partition, the partition must be either in the Not activated or the Running state. If the partition is in the Running state, the partition must also be capable of dynamic partitioning (DLPAR).

To avoid configuring the physical Fibre Channel adapter to be a single point of failure for the connection between the client logical partition and its physical storage on the SAN, do not connect two virtual Fibre Channel adapters from the same client logical partition to the same physical Fibre Channel adapter. Instead, connect each virtual Fibre Channel adapter to a different physical Fibre Channel adapter.

To change virtual Fibre Channel, be sure you meet the following requirements:

- The HMC must be at version 7 release 3.4.2 or later.
- The Virtual I/O Server must be at version 2.1.1.0 or later.

- Ensure that there is a resource monitoring and control connection between the HMC and the Virtual I/O Server.

About this task

To configure the physical port connections for virtual Fibre Channel, complete the following steps in the HMC:

Procedure



1. In the navigation pane, click the **Resources** icon.
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the **PowerVM** area, click **Virtual Storage**. The Virtual Storage page opens with the VIOS partitions listed in a table, in the Virtual Storage Management tab.
4. Select a VIOS and click **Action > Manage Virtual Storage**.
5. Select a Virtual I/O Server logical partition.
6. Click the **Virtual Fibre Channel** tab.
7. Select a port with at least one available connection and click **Modify partition connections**.
The Modify Virtual Fibre Channel Partition Assignment page is displayed.
8. Select one or more logical partitions that you want to connect to the Fibre Channel port, and click **OK**.

Note: If you delete the client virtual Fibre Channel adapter from the partition or the partition profile, the worldwide port names associated with the port and the storage area network (SAN) are lost. If you only change the port assignment, the worldwide port names are preserved inside the partition profile. The HMC does not reuse them when it generates port names in the future. If you run out of port names, you must obtain a code key to enable an additional prefix and range of port names for use on your system.

9. Click **OK**.

To determine the actual number of port names available on the managed system, use the HMC to view the partition properties or partition profile properties of the client logical partition.

Managing the memory configuration of a logical partition

You can use the Hardware Management Console (HMC) to change the memory configuration of a logical partition. For example, you can change the Virtual I/O Server logical partitions that are assigned to a logical partition that uses shared memory, change the memory mode of a logical partition, and dynamically add and remove dedicated or shared memory to and from a logical partition.

Changing the paging VIOS partitions assigned to a shared memory partition

You can use the Hardware Management Console (HMC) to change the primary and secondary Virtual I/O Server logical partitions (hereafter referred to as *paging VIOS partitions*) that are assigned to a logical partition that uses shared memory. You can also add or remove a secondary paging VIOS partition to or from a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*).

Before you begin

Before you change the paging VIOS partitions that are assigned to a shared memory partition, complete the following tasks:

1. Ensure that the Virtual I/O Server logical partitions (that you plan to assign to the shared memory partition as paging VIOS partitions) are assigned to the shared memory pool. For instructions, see [“Changing the paging VIOS partitions assigned to the shared memory pool” on page 127](#).
2. Ensure that the paging space device (that is accessed through the paging VIOS partitions that you plan to assign to the shared memory partition) is assigned to the shared memory pool. For instructions, see [“Adding and removing paging space devices to and from the shared memory pool” on page 132](#).

About this task

To change the paging VIOS partitions that are assigned to a shared memory partition, complete the following steps:

Procedure



1. In the navigation pane, click the **Resources** icon.
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select the logical partition and click **Actions > Profiles > Manage Profiles**.
4. Select the partition profile that you want to change.
5. Click **Actions**, and click **Edit**.
The Logical Partition Profile Properties window is displayed.
6. Click the **Memory** tab.
7. Specify a Virtual I/O Server logical partition for VIOS 1 and VIOS 2.

Table 25. Change options for the paging VIOS partitions	
Desired change	Field to change
Change the Virtual I/O Server logical partition that is assigned as the primary or the only paging VIOS partition.	Select a different Virtual I/O Server logical partition for VIOS 1.
Define a secondary paging VIOS partition.	Select a Virtual I/O Server logical partition for VIOS 2.
Change the Virtual I/O Server logical partition that is assigned as the secondary paging VIOS partition.	Select a different Virtual I/O Server logical partition for VIOS 2.
Remove the secondary paging VIOS partition.	Select None for VIOS 2.

8. Click **OK**.
9. Shut down the shared memory partition and reactivate it with the changed partition profile.

What to do next

After you change the paging VIOS partitions that are assigned to a shared memory partition, restart the shared memory partition with the changed partition profile. For instructions, see [“Shutting down and restarting logical partitions” on page 137](#).

Changing the Active Memory Expansion factor for AIX logical partitions

You can dynamically change the Active Memory Expansion factor for an AIX logical partition by using the Hardware Management Console (HMC). Changing the Active Memory Expansion factor for a logical partition increases or decreases the desired degree of expanded memory capacity for the logical partition.

Before you begin

You can change the Active Memory Expansion factor for logical partitions that use dedicated memory and logical partitions that use shared memory.

About this task

For more information about changing the memory settings, see [Changing memory settings](#).

Changing the memory weight of a shared memory partition

You can use the Hardware Management Console (HMC) to change the memory weight of a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*). Changing the memory weight changes the probability that the shared memory partition receives physical memory from the shared memory pool in relation to other shared memory partitions.

Before you begin

A Linux shared memory partition supports changing the memory weight only if the DynamicRM tool package is installed on the Linux shared memory partition. To download the DynamicRM tool package, see the [Service and productivity tools for Linux on POWER systems website](#).

About this task

For more information about changing memory settings on a logical partition, see [Changing memory settings](#).

What to do next

Changing the memory weight of a shared memory partition is temporary and is not reflected in the partition profile. The new memory weight that you assigned to the shared memory partition will be lost the next time you activate the partition profile. If you want to save the changes that you made to the memory weight of the shared memory partition, either change the partition profile or save the logical partition configuration to a new partition profile.

Changing the memory mode of a logical partition

You can create multiple partition profiles for a logical partition by using the Hardware Management Console (HMC). Some of the partition profiles can specify dedicated memory and some of the partition profiles can specify shared memory. By creating partition profiles that specify both dedicated memory and shared memory for the same logical partition, you can change the memory mode of the logical partition by activating different partition profiles.

About this task

To change the memory mode of a logical partition, complete the following steps from the HMC:

Procedure

1. Create a new partition profile for the logical partition.
For instructions, see [“Creating additional partition profiles” on page 92](#).
 - If you plan to change a dedicated memory partition to a shared memory partition, specify the shared memory mode in the new partition profile.
 - If you plan to change a shared memory partition to a dedicated memory partition, specify the dedicated memory mode in the new partition profile.
2. Shut down the logical partition.
For instructions, see [“Shutting down and restarting logical partitions” on page 137](#).
3. Activate the logical partition with the new partition profile.
For instructions, see [“Activating a partition profile” on page 133](#).

Managing dedicated memory dynamically

You can add, remove, and move physical memory dynamically to and from running logical partitions that use dedicated memory by using the Hardware Management Console (HMC). This allows you to adjust the

physical memory allocated to each logical partition that uses dedicated memory without having to shut down the logical partitions.

When a DPO operation is in progress and you want to dynamically add, remove, or move physical memory to, or from running logical partitions, you must either wait for the DPO operation to complete, or manually stop the DPO operation.

Dynamic memory changes on IBM i logical partitions affect the base memory pool of the logical partitions (*BASE pool). Private memory pools or shared memory pools are not affected. Dynamic memory changes cannot cause the amount of memory in the base pool to fall below the minimum amount of memory required in the base pool (as determined by the base storage minimum size (QBASPOOL) system value). If a dynamic memory change would cause the base pool to fall below this amount, the system releases excess memory pages only after keeping the minimum amount of memory required in the base pool.

To prevent any data loss during dynamic memory movement, the system first writes any data from memory pages to disk before making the memory pages available to another logical partition. Depending on the amount of memory you have requested to move, this might take some time.

Memory in each logical partition operates within its assigned minimum and maximum values. The full amount of memory that you assign to a logical partition might not be available for the logical partition to use. Static memory overhead that is required to support the assigned maximum memory affects the reserved or hidden memory amount. This static memory overhead also influences the minimum memory size of a logical partition.

Note:

- If resources are moved dynamically, the configuration change is temporary and is not reflected in the partition profile. This means that all configuration changes will be lost the next time the partition profile is activated. If you want to save your new logical partition configuration, either change the partition profile or save the logical partition configuration to a new partition profile.
- When dynamic logical partitioning tasks to add, remove, or move physical memory are run concurrently for a logical partition, the logical partition might not have the expected amount of physical memory after the concurrent tasks are complete. The logical partition might not have the expected amount of physical memory whether you specify the amount of physical memory that you want the logical partition to have after the dynamic logical partitioning task is complete, or you specify the amount of physical memory to be added to, removed from, or moved to or from the logical partition.

Adding dedicated memory dynamically

You can dynamically add physical memory to a running logical partition that uses dedicated memory using the Hardware Management Console (HMC). This allows you to increase the physical memory available to a logical partition that uses dedicated memory without having to shut down the logical partition.

Before you begin

A Linux logical partition supports the dynamic addition of memory resources only if the following conditions are met:

- A Linux distribution that supports the dynamic addition of memory resources is installed on the Linux logical partition. Distributions that support the dynamic addition of memory resources include SUSE Linux Enterprise Server 10, and later.
- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the [Service and productivity tools for Linux on POWER systems](#) website.

To add memory to a Linux logical partition that uses an earlier version of these distributions, you must shut down the Linux logical partition and reactivate the logical partition using a partition profile that specifies a greater amount of memory.

About this task

For more information about changing the memory settings, see [Changing memory settings](#).

Moving dedicated memory dynamically

You can dynamically move physical memory from one running logical partition that uses dedicated memory to another using the Hardware Management Console (HMC). This allows you to reassign physical memory directly to a logical partition that uses dedicated memory that needs additional physical memory.

Before you begin

You cannot dynamically move memory from a running Linux logical partition. To remove memory from a Linux logical partition, you must shut down the Linux logical partition and reactivate the logical partition using a partition profile that specifies a lesser amount of memory.

You can dynamically move memory to a running Linux only if the following conditions are met:

- A Linux distribution that supports the dynamic addition of memory resources is installed on the Linux logical partition. Distributions that support the dynamic movement of memory resources include Novell SUSE Linux Enterprise Server 10, and later.
- The DynamicRM tool package is installed on the Linux logical partition. To download the DynamicRM tool package, see the [Service and productivity tools for Linux on POWER systems](#) website.

To move memory to a Linux logical partition that uses an earlier version of these distributions, you must shut down the Linux logical partition and reactivate the logical partition using a partition profile that specifies a greater amount of memory.

About this task

For more information about changing the memory settings, see [Changing memory settings](#).

Removing dedicated memory dynamically

You can dynamically remove physical memory from a running AIX, IBM i, or Virtual I/O Server logical partition that uses dedicated memory using the Hardware Management Console (HMC). This allows you to reassign the physical memory to other logical partitions that use dedicated memory.

Before you begin

You cannot dynamically remove memory from a running Linux logical partition. To remove memory from a Linux logical partition, you must shut down the logical partition and reactivate the logical partition using a partition profile that specifies a lesser amount of memory.

About this task

For more information about changing the memory settings, see [Changing memory settings](#).

Managing shared memory dynamically

You can dynamically add and remove logical memory and I/O entitled memory to and from a logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) using the Hardware Management Console (HMC).

About this task

Dynamic memory changes on IBM i logical partitions affect the base memory pool of the logical partitions (*BASE pool). Private memory pools or shared memory pools are not affected. Dynamic memory changes cannot cause the amount of memory in the base pool to fall below the minimum amount of memory required in the base pool (as determined by the base storage minimum size (QBASPOOL) system value). If a dynamic memory change would cause the base pool to fall below this amount, the system releases excess memory pages only after keeping the minimum amount of memory required in the base pool.

Note:

- If resources are moved dynamically, the configuration change is temporary and is not reflected in the partition profile. This means that all configuration changes will be lost the next time the partition profile

is activated. If you want to save your new logical partition configuration, either change the partition profile or save the logical partition configuration to a new partition profile.

- When dynamic logical partitioning tasks to add or remove shared memory are run concurrently for a logical partition, the logical partition might not have the expected amount of shared memory after the concurrent tasks are complete. The logical partition might not have the expected amount of shared memory whether you specify the amount of shared memory that you want the logical partition to have after the dynamic logical partitioning task is complete, or you specify the amount of shared memory to be added to or removed from the logical partition.

Adding and removing logical memory dynamically to and from a shared memory partition

You can dynamically add and remove logical memory to and from a running logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) using the Hardware Management Console (HMC). This allows you to increase and decrease the logical memory assigned to the shared memory partition without having to shut down the logical partition.

Before you begin

A Linux shared memory partition supports the dynamic addition and removal of logical memory resources only if the DynamicRM tool package is installed on the Linux shared memory partition. To download the DynamicRM tool package, see the [Service and productivity tools for Linux on POWER systems](#) website.

To dynamically add and remove logical memory to and from a running logical partition using the HMC, you must be a super administrator, service representative, product engineer, or operator of the HMC.

About this task

For more information about changing the memory settings, see [Changing memory settings](#).

Adding and removing I/O entitled memory dynamically to and from a shared memory partition

You can dynamically add and remove I/O entitled memory to and from a running logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) using the Hardware Management Console (HMC). This allows you to increase and decrease the maximum amount of physical memory that is assigned to the shared memory partition for its I/O devices without having to shut down the shared memory partition.

Before you begin

A Linux shared memory partition supports the dynamic addition and removal of I/O entitled memory resources only if the DynamicRM tool package is installed on the Linux shared memory partition. To download the DynamicRM tool package, see the [Service and productivity tools for Linux on POWER systems](#) website.

You can increase the amount of I/O entitled memory that is assigned to a shared memory partition when the sum of I/O entitled memory that is assigned to all shared memory partitions in the shared memory pool is less than the size of the shared memory pool minus the required amount of reserved firmware memory. If there is not enough physical memory in the shared memory pool by which to increase the I/O entitled memory to the amount specified, you can release to the hypervisor the physical memory that is currently assigned to other shared memory partitions that are shut down. The hypervisor can then assign the released physical memory to the shared memory partition that needs more I/O entitled memory.

You can decrease the amount of I/O entitled memory that is assigned to a shared memory partition only when the shared memory partition requires less physical memory for its I/O devices than the amount of I/O entitled memory that is assigned to the shared memory partition. For example, you assign 128 MB of I/O entitled memory to a shared memory partition. The shared memory partition requires a minimum of 64 MB for its I/O devices. Thus, you can decrease the I/O entitled memory that is assigned to the shared memory partition by up to 64 MB. For instructions about how to view the assigned, minimum, optimal, and maximum I/O entitled memory used by a shared memory partition, see [“Determining the I/O entitled memory for a shared memory partition”](#) on page 219.

To dynamically add and remove I/O entitled memory to and from a running shared memory partition using the HMC, you must be a super administrator, service representative, product engineer, or operator of the HMC.

About this task

For more information about changing the memory settings, see [Changing memory settings](#).

Results

If you want to later change the I/O entitled memory mode back to the auto mode so that the HMC automatically adjusts the I/O entitled memory for the shared memory partition when you add or remove virtual adapters, repeat this procedure and select **Auto**. Alternatively, you can restart the shared memory partition. When you restart a shared memory partition, the I/O entitled memory mode is set to the auto mode regardless of what the I/O entitled memory mode was set to before you restarted the shared memory partition.

Obtaining additional WWPNS for the server

When all of the worldwide port names (WWPNs) on the server are used, you can add more WWPNS to the server using the Hardware Management Console (HMC). Adding WWPNS allows you to create additional virtual Fibre Channel adapters on client logical partitions that use virtual resources provided by the Virtual I/O Server.

Before you begin

The server contains 32,000 pairs of WWPNS that all contain the same 6–digit prefix. Each virtual Fibre Channel adapter that you create on a client logical partition requires one pair of WWPNS. When all of the WWPNS on the server are used, you cannot create additional virtual Fibre Channel adapters on any client logical partitions until you add more WWPNS to the server. You add more WWPNS to the server by generating an activation code that contains a new WWPN prefix that contains 32,000 new pairs of WWPNS.

About this task

To obtain additional WWPNS for the server, complete the following steps from the HMC:

Procedure

1. Retrieve information about the server:



- a) In the navigation pane, click the **Resources** icon.
 - b) Click **All Systems**. The **All Systems** page is displayed.
 - c) In the work pane, select the system and click **Actions > View System Properties**. The **Properties** page is displayed.
The CoD Advanced Functions Code Information window is displayed.
 - d) In the **CUoD (permanent) Processor** area, click **View CUoD Code Information**.
 - e) Click **Save** to save the information to a file on a remote system or to media, and click **OK**.
2. Go to the Capacity on Demand website and enter the information that you retrieved in step “1” on [page 180](#) to generate an activation code.
 3. Apply the activation code that you obtained in step “2” on [page 180](#) to the server:



- a) In the navigation pane, click the **Resources** icon.
- b) Click **All Systems**. The **All Systems** page is displayed.

- c) In the work pane, select the system and click **Actions > View System Properties**. The **Properties** page is displayed.
The CoD Advanced Functions Code Information window is displayed.
- d) In the **Capacity on Demand** area, click **CoD Functions**.
- e) Enter the activation code that obtained in step “2” on page 180 and click **OK**.
- 4. Verify that the activation code that you entered in step “3” on page 180 was applied to the server:



- a) In the navigation pane, click the **Resources** icon .
- b) Click **All Systems**. The **All Systems** page is displayed.
- c) In the work pane, select the system and click **Actions > View System Properties**. The **Properties** page is displayed.
The CoD Advanced Functions Code Information window is displayed.
- d) In the **Capacity on Demand** area, click **CoD Functions**.
- e) In the Capacity On Demand Functions page, click **View CoD History Log**.
- f) Verify that there is a log entry for entering the CoD advanced functions activation code and click **Close**.

What to do next

After you finish, you can create virtual Fibre Channel adapters on client logical partitions and dynamically add virtual Fibre Channel adapters to client logical partitions.

Related concepts

Virtual Fibre Channel

With N_Port ID Virtualization (NPIV), you can configure the managed system so that multiple logical partitions can access independent physical storage through the same physical Fibre Channel adapter.

Setting partition-availability priorities for your managed system

To avoid shutting down mission-critical workloads when your server firmware deconfigures a failing processor, you can use the Hardware Management Console (HMC) to set partition-availability priorities for the logical partitions on your managed system. A logical partition with a failing processor can acquire a replacement processor from logical partitions with a lower partition-availability priority. The acquisition of a replacement processor allows the logical partition with the higher partition-availability priority to continue running after a processor failure.

About this task

To set partition-availability priorities for your managed system by using the HMC, follow these steps:

Procedure



1. In the navigation pane, click the **Resources** icon .
2. Click **All Systems**. The **All Systems** page is displayed.
3. In the work pane, select the system and click **System Actions > Legacy > Partition Availability Priority**.
4. Select the logical partitions whose partition-availability priority you want to set, set **Availability priority** to the partition-availability priority value that you want to use for all selected logical partitions, and click **OK**.

You can enter any value from 0 to 255 into **Availability priority**, or you can select one of the preset choices. All selected logical partitions are set to the same partition-availability priority value.

5. Repeat this procedure for other logical partitions to set the partition-availability priority for those logical partitions.

Installing new hardware for IBM i logical partitions

You can install an I/O adapter (IOA) for an IBM i logical partition.

Before you begin

When you install new hardware in an IBM i partitioned environment, you should be aware of the following things:

- Verify that your logical partition configuration is current.
- Empty positions might not be owned by a logical partition. They should be assigned to the desired logical partition before installing new adapters in them. After you install the new adapter, you must also add the adapter to the partition profile so that, when you shut down and activate the logical partition using the partition profile, the logical partition reactivates with the adapter that you added.
- A new IOA is owned by the logical partition that owns the slot, and a new device is owned by the logical partition that owns the IOA to which the device is attached.
- New processors and memory are available (unassigned) to be assigned to any logical partition.

About this task

To install an IOA for an IBM i logical partition, perform the following steps:

Procedure

1. Assign empty slots to the desired logical partition.
For instructions, see [“Managing physical I/O devices and slots dynamically”](#) on page 161 and [“Changing partition profile properties”](#) on page 149.
2. Install the new hardware into the empty slots. For instructions, see [Installing and configuring POWER9 processor-based systems and system features](#).

Backing up and recovering data

It is crucial that you back up your data because you never know when you might need to do a server recovery. Save everything in your system as often as possible. You might not be prepared to recover from a site loss or certain types of disk failures if you do not regularly save everything.

For more information about planning a backup and recovery strategy for the Hardware Management Console (HMC) and IBM i data, refer to the following topics:

Table 26. Backup and recovery information for the HMC and IBM i and IBM i	
Topic	Description
Backing up critical HMC data	This procedure explains how to save critical HMC data (such as user information and platform-configuration files) to a backup file. This information is in the Managing the HMC topic.
Backing up partition profile data	This procedure explains how to back up the partitioning data on your HMC to a backup file on the HMC. This information is in the Managing the HMC topic.
Reinstalling the HMC machine code	This procedure explains how to reinstall the HMC interface from the recovery CD-ROM. This information is in the Managing the HMC topic.

Table 26. Backup and recovery information for the HMC and IBM i and IBM i (continued)	
Topic	Description
Restoring profile data	This procedure explains how to restore the partitioning data from the backup file to the HMC. This information is in the Managing the HMC topic.
Back up your server	This information can help you develop the backup strategy for your IBM i logical partition. This information is in the Backup and recovery topic in the IBM i Knowledge Center.
Recover your server	This information can help you reload your operating system and data. This information is in the Backup and recovery topic in the IBM i Knowledge Center.

Managing logical partitions that use IBM i resources

You can manage logical partitions that use IBM i virtual I/O resources to help maximize utilization of the physical hardware and simplify the backup procedure for your managed system.

Managing AIX logical partitions that use IBM i resources

You can manage AIX logical partitions that uses IBM i virtual I/O resources to help maximize utilization of the physical hardware and simplify the backup procedure for your managed system.

Related information

[Backup of the system image and user-defined volume groups](#)

[Installing system backups](#)

Adding virtual disk units to an AIX logical partition

You can dynamically add virtual disk units to an AIX logical partition that uses IBM i resources. This allows you to increase the storage capacity of your AIX logical partition when needed.

About this task

Virtual disks simplify hardware configuration on the server because they do not require you to add additional physical devices to the server in order to run AIX. You can allocate up to 64 virtual disks to an AIX logical partition. Each virtual disk supports up to 1000 GB of storage. Each virtual disk appears to AIX as one actual disk unit. However, the associated space in the IBM i integrated file system is distributed across the disks that belong to the IBM i logical partition. Distributing storage across the disks provides the benefits of device parity protection through IBM i. Therefore, you do not have to use additional processing resources and memory resources by setting up device parity protection through AIX.

IBM i provides the ability to dynamically add virtual disks to an AIX logical partition. You can allocate disk space in the integrated file system and make it available to AIX without restarting the server or logical partition. The AIX administrator can also configure the newly allocated disk space and make it available without restarting the server.

To add virtual disks dynamically to an AIX logical partition, complete the following steps:

Procedure

1. If you use IBM Navigator for i, create a network-server storage space using IBM Navigator for i.
 - a) Expand **My Connections** > **your server** > **Network** > **Windows Administration** .
 - b) Right-click the **Disk Drives** and select **New Disk**.
 - c) In the **Disk drive name** field, specify the name that you want to give to the network-server storage space.

- d) In the **Description** field, specify a meaningful description for the network-server storage space.
 - e) In the **Capacity** field, specify the size of the new network-server storage space in megabytes.
To help you determine the size you want to use, see [Installing AIX](#).
 - f) Click **OK**.
 - g) Continue with step “3” on page 184.
2. If you use a character-based interface, create a network-server storage space using the character-based interface:
 - a) At an IBM i command line, type the command **CRTNWSSTG** and press F4.
The Create NWS Storage Space (**CRTNWSSTG**) display opens.
 - b) In the Network-server storage space field, specify the name you want to give to the network-server storage space.
 - c) In the Size field, specify the size in megabytes for the new network-server storage space.
To help you determine the size you want to use, see [Installing AIX](#).
 - d) In the Text description field, specify a meaningful description for the network-server storage space.
 - e) Press Enter.
 - f) Continue with step “4” on page 184.
 3. If you use IBM Navigator for i, link the network-server storage space using IBM Navigator for i.
 - a) Expand **My Connections > your server > Network > Windows Administration**.
 - b) Click **Disk Drives**, right-click an available network-server storage space, and select **Add Link**.
 - c) Select the server to which you want to link the network-server storage space.
 - d) Select one of the available data access types.
 - e) Click **OK**.
 - f) Continue with step “5” on page 184.
 4. If you use a character-based interface, link the network-server storage space using a character-based interface:
 - a) At an IBM i command line, type the command **ADDNWSSTGL** and press F4.
The Add Network-Server Storage Link (**ADDNWSSTGL**) display opens.
 - b) In the Network server description field, specify the name of the network server description (NWSD).
 - c) In the Dynamic storage link field, specify *YES to make the network-server storage space dynamically available to the logical partition (that is, available without rebooting the AIX logical partition).
 - d) In the Drive sequence number field, specify the link sequence position you want to use.
 - e) Press Enter.
 5. Activate the AIX logical partition (if it is not already activated).
 6. Log in to AIX using a user name with superuser (root) privileges.
 7. Configure the new virtual disk on the AIX logical partition by running the AIX command `cfgmgr`.
 8. Verify that your new disk has been added and can be configured by running the AIX command `lspv`.
When you enter `lspv` at the command prompt, the system lists the disks that are currently available to AIX.

An example of the output for this command is below:

# lspv			
hdisk0	00cad6aceafe8fe4	rootvg	active
hdisk1	none	None	

Note the name of the new disk as it displays in the left-hand column.

9. Configure the new disk using one of the following two methods.

- Add the new virtual disk to the root volume group by using the AIX command `extendvg rootvg diskname`, where *diskname* is the name of the new disk. If you use this method, you do not need to continue this procedure. You can use AIX methods to increase the file system size at a later time.
- Create a new volume group for the new virtual disk by using the AIX command `mkvg -y volgroup diskname`, where *volgroup* is the name that you want to use for the new volume group and *diskname* is the name of the new disk.

10. Make a logical volume on the new virtual disk using the AIX `mklv -y logicvol volgroup 1 diskname` command.

logicvol is the name that you want to use for the new logical volume, *volgroup* is the name of the new volume group, and *diskname* is the name of the new disk. (The numeral 1 indicates that the logical volume is to consist of one logical disk partition.)

11. Format the disk partition using the AIX `crfs` command.

There are a number of optional parameters for the `crfs` command, but typically the defaults satisfy most disk uses. To format the disk partition created in the previous steps, type the following command at an AIX command prompt, where *logicvol* is the name of the logical volume and */mnt/data* is the mount point directory at which you want to mount the new disk:

```
crfs -v jfs -d logicvol -m /mnt/data
```

The `crfs` command displays the following diagnostic messages:

```
crfs -v jfs -d logicvol -m /mnt/data
Based on the parameters chosen, the new /mnt/data JFS file system is limited to
a maximum size of 134217728 (512 byte blocks)
New File System size is 8192.
```

12. Verify that the mount point directory exists by using the `cd /mnt/data` command.

/mnt/data is the mount point. The `crfs` command creates this directory so that you can access your new file system. If the mount point directory does not exist, then run the following command, where */mnt/data* is the name of the mount point directory:

```
mkdir /mnt/data
```

13. Verify that an entry for your new file system exists in the */etc/filesystems* file.

The `crfs` command automatically generates the appropriate */etc/filesystems* entry for your new file system. To verify that the entry exists, use an AIX text editor, such as `vi`, to open the */etc/filesystems* file, and look for the entry in the */etc/filesystems* file. If the entry does not exist, use the text editor to add the entry to the */etc/filesystems* file.

An example of such an entry is below:

```
/mnt/data:
dev = /dev/logicvol
vfs = jfs
log = /dev/loglv01
mount = true
account = false
```

This entry mounts the virtual disk every time you restart AIX.

14. Mount the virtual disk drive in the new directory by typing: `mount /dev/logicvol /mnt/data`.

logicvol is the name of the logical volume and */mnt/data* is the mount point directory.

Linking a network-server storage space to a network server description

You can link a network-server storage space (NWSSTG) to one or more network server descriptions (NWSDs). This allows the NWSDs and their associated logical partitions to use the data stored on the NWSSTG.

About this task

You can link an NWSSTG to an unlimited number of NWSDs. This is beneficial when multiple logical partitions need access to a single application.

When you link an NWSSTG to an NWSD, you can set up the NWSD to have read-only access to the NWSSTG, or you can set up the NWSD to read or write to the NWSSTG.



Attention: If more than one NWSD can write to the NWSSTG, ensure that only one NWSD can update the data at a time. Otherwise, changes made by one NWSD can be overwritten by another NWSD.

To link an NWSSTG to an NWSD, follow these steps:

Procedure

1. At an IBM i command line, type the command ADDNWSSTGL and press F4.
2. From the Add Server Storage Link display, provide the following information:

```
NWSSTG (Name)
NWSD (Name)
DYNAMIC (*YES)
DRVSEQNBR (*CALC)
```

3. Press F10 (Additional Parameters).
4. Enter the type of access the storage space will have.

Deleting network-server descriptions for an AIX logical partition

You can delete the IBM i network-server description (NWSD) for an AIX logical partition that uses IBM i resources. When you delete the NWSD, all the configuration information for the AIX logical partition is deleted from IBM i.

About this task

To delete the network-server description (NWSD) for an AIX logical partition, follow these steps:

Procedure

1. On an IBM i control language (CL) command line, type the command WRKNWSD and press Enter.
2. Type 8 in the Opt field to the left of the Network Server and press Enter.
3. In the Work with Configuration Status display, if the status of the NWSD is not varied off, type 2 in the Opt field to the left of the Network Server and press Enter. Otherwise, go to the next step.
4. Press F3 to return to the previous display
5. Enter a 4 in the Opt field to the left of the Network Server and press Enter.
6. On the Confirm Delete of Network Server Descriptions display, press Enter.

Deleting virtual disk drives for an AIX logical partition

You can delete a virtual disk drive from an AIX logical partition that uses IBM i resources to make the space available to the IBM i logical partition once more. When you delete a virtual disk drive, all of the information on the virtual disk drive is erased.

Before you begin

Before you can delete a virtual disk drive, you must unlink the virtual disk drive from the network-server description (NWSD). For instructions, see [“Unlinking virtual disk drives from an AIX logical partition” on page 188](#).

About this task

To delete a virtual disk drive, follow these steps:

Procedure

Delete the disk drive using the interface that you prefer.

Interface	Actions
IBM Navigator for i	- Click Network > Windows Administration > Disk Drives. - Right-click the disk drive that you want to delete. - Click Delete in the confirmation window.
IBM i character-based interface	- At an IBM i control language (CL) command line, type DLTNWSSTG and press F4. - Type the name of the disk drive in the Network-server storage space field and press Enter.

Using IPL types when running AIX

The IPL source (IPLSRC) parameter on the network-server description (NWSD) determines the initial program that is loaded when the NWSD is varied on. For an AIX logical partition that uses IBM i resources, the initial program is the kernel. Ensure that the IPLSRC parameter specifies the kernel location of the kernel for the AIX logical partition that uses IBM i resources.

You can set the IPLSRC parameter when you use the Create Network Server Description (CRTNWSD) command, and you can change the IPLSRC parameter when you use the Change Network Server Description (CHGNWSD) command.

Note: The IPLSRC parameter also has the values A, B, and D, which are not valid for hardware that is used by IBM i logical partitions.

The IPLSRC parameter has the following valid values.

IPLSRC values	Description
*Panel	The logical partition is started from the source indicated on the control panel.
*NWSSTG (network-server storage space)	This IPL type is used to start a logical partition from a virtual disk. The open firmware will find the kernel in the virtual disk. The open firmware searches the first virtual disk connected to the server for a logical partition marked bootable, and of type 0x41 (PReP start). If a logical partition of this type does not exist, the logical partition IPL will fail.
*STMF (stream file)	This IPL type is used to start a logical partition from a kernel IBM i loaded in the IBM i integrated file system. Note that the integrated file system includes files on the optical (CD) drive on IBM i.

Unlinking virtual disk drives from an AIX logical partition

By unlinking virtual disk drives (network-server storage spaces) from an AIX logical partition that uses IBM i resources, you disconnect the virtual disk drives from the logical partition, making the virtual disk drives inaccessible to users. If you delete an AIX logical partition that uses IBM i resources, you must unlink all virtual disk drives from the logical partition before you delete the logical partition.

About this task

To unlink a virtual disk drive from an AIX logical partition that uses IBM i resources, follow these steps:

Procedure

1. Unlink disk drives from a logical partition by using IBM Navigator for i.
If you prefer to use a character-based interface, go to step [“2” on page 188](#).
 - a) Vary off the NWSD for your logical partition.
 - b) Click **Network > Windows Administration > Disk Drives**.
 - c) Right-click the name of the disk drive that you want to unlink.
 - d) Click **Remove Link**.
 - e) Select a server from the list of linked servers.
 - f) If you are unlinking a disk drive that you plan to relink later, clear **Compress link sequence**.
You must relink the disk drive as the same link sequence number before you vary on the server. By preventing compression of the link sequence values, you avoid having to unlink and relink all the disk drives to get them in the correct sequence.
 - g) Click **Remove**.
 - h) You have completed this procedure. Do not complete step [“2” on page 188](#).
2. Unlink disk drives from a logical partition that uses a character-based interface:
 - a) Vary off the NWSD for your logical partition.
 - b) Type RMVNWSTGL and press F4.
 - c) In the **Network-server storage space** field, type the name of the storage space that you want to unlink and press Enter.
 - d) In the **Network server description** field, type the name of the server from which you want to unlink the storage space and press Enter.
 - e) If you are unlinking a linked disk drive that you plan to relink later, specify *NO in the **Renumber** field.
Note: You must relink the disk drive as the same sequence number before you vary on the server. By preventing automatic renumbering, you avoid having to unlink and relink all the disk drives to get them in the correct sequence.
 - f) Press Enter.
Note: If you are uninstalling a logical partition, your next step is to delete the disk drive. For instructions, see [“Deleting virtual disk drives for an AIX logical partition” on page 187](#). Otherwise, vary on the NWSD for your logical partition.

Saving AIX server objects in IBM i

When an AIX logical partition uses IBM i resources, IBM i stores AIX information in IBM i objects. IBM i can restore the objects correctly only if you save all objects for an AIX logical partition.

You can save these objects by using options of the IBM i GO SAVE command in the server.

- Option 21 saves the entire server.
- Option 22 saves server data, which includes objects in the QUSRSYS library.
- Option 23 saves all user data, which includes objects in the QFPNWSSTG library.

If you want to save a particular object, use the following table to see the location of that object on IBM i and the command to use.

<i>Table 27. Objects to save for logical partitions with virtual disk</i>				
Object content	Object name	Object location	Object type	Save command
Guest partition and virtual disk drive	stgspc	/QFPNWSSTG	User-defined network-server storage spaces in system auxiliary storage pool (ASP)	GO SAV, option 21 or 23
				SAV OBJ('/QFPNWSSTG/stgspc') DEV('/QSYS.LIB/TAP01.DEVD')
			User-defined network-server storage spaces in user ASP	SAV OBJ('/QFPNWSSTG/stgspc') ('/dev/QASPnn /stgspc.UDFS')) DEV('/QSYS.LIB/TAP01.DEVD')

<i>Table 28. Objects to save for all logical partitions with a server</i>				
Object content	Object name	Object location	Object type	Save command
Messages from the logical partition	Various	Various	Server message queue	GO SAVE, option 21 or 23
				SAVOBJ OBJ(msg) LIB(qlibrary) DEV(TAP01) OBJTYPE(*MSGQ)
IBM i configuration objects for logical partitions	Various	QSYS	Device configuration objects	GO SAVE, option 21, 22, or 23
				SAVOBJ DEV (TAP01)
Various	Various	QUSRSYS	Various	GO SAVE, option 21 or 23
				SAVLIB LIB(*NONSYS) or LIB(*ALLUSR)

Related information

[Backup of the system image and user-defined volume groups](#)

[Installing system backups](#)

Managing IBM i logical partitions that use i resources

You can manage IBM i logical partitions that uses i virtual I/O resources to help maximize utilization of the physical hardware and simplify the backup procedure for your managed system.

Adding virtual disk units to an IBM i logical partition that uses i virtual I/O resources

You dynamically can add virtual disk units to an IBM i logical partition that uses i virtual I/O resources. This allows you to increase the storage capacity of the i logical partition that uses i virtual I/O resources when needed.

About this task

IBM i provides the ability to add virtual disks dynamically to another i logical partition. You can allocate disk space in the integrated file system and make it available to i without restarting the server or logical partition.

To add virtual disks dynamically to an IBM i logical partition that uses i virtual I/O resources, complete the following steps:

Procedure

1. Create a network-server storage space using the interface of your choice.

Interface	Actions
IBM Navigator for i	a. Expand My Connections > your server > Network > Windows Administration. b. Right-click the Disk Drives and select New Disk. c. In the Disk drive name field, specify the name that you want to give to the network-server storage space. d. In the Description field, specify a meaningful description for the network-server storage space. e. In the Capacity field, specify the size of the new network-server storage space in megabytes. To help you determine the size you want to use, see Installing, upgrading, or deleting IBM i and related software. f. Click OK.
IBM i character-based interface	a. At an IBM i command line on the IBM i logical partition that provides virtual I/O resources, type the command CRTNWSSTG and press F4. The Create NWS Storage Space (CRTNWSSTG) display opens. b. In the Network-server storage space field, specify the name you want to give to the network-server storage space. c. In the Size field, specify the size in megabytes for the new network-server storage space. d. In the Text description field, specify a meaningful description for the network-server storage space. e. Press Enter.

2. Add the new disk to the auxiliary storage pool (ASP) on the client IBM i logical partition.
For instructions, see [Adding disk units to an existing auxiliary storage pool \(ASP\)](#).
3. Link the network-server storage space using the interface of your choice.

Interface	Actions
IBM Navigator for i	a. Expand My Connections > your server > Network > Windows Administration. b. Click Disk Drives, right-click an available network-server storage space, and select Add Link. c. Select the server to which you want to link the network-server storage space. d. Select one of the available data access types. e. Click OK.
IBM i character-based interface	a. At an IBM i command line on the IBM i logical partition that provides virtual I/O resources, type the command ADDNWSSTGL and press F4. The Add Network-Server Storage Link (ADDNWSSTGL) display opens. b. In the Network server description field, specify the name of the network server description (NWSD). c. In the Dynamic storage link field, specify *YES to make the network-server storage space dynamically available to the logical partition (that is, available without restarting the IBM i logical partition). d. In the Drive sequence number field, specify the link sequence position you want to use. e. Press Enter.

4. Activate the IBM i logical partition (if it is not already activated).

Linking a network-server storage space to a network server description

You can link a network-server storage space (NWSSTG) to a network server description (NWSD). This allows the NWSD and its associated logical partition to use the data stored on the NWSSTG.

About this task

You can link an NWSSTG to one NWSD. When you link an NWSSTG to an NWSD, you can set up the NWSD to have read-only access to the NWSSTG, or you can set up the NWSD to read or write to the NWSSTG.

To link an NWSSTG to an NWSD, follow these steps:

Procedure

1. At an IBM i command line on the IBM i logical partition that provides virtual I/O resources, type the command ADDNWSSTGL and press F4.
2. From the Add Server Storage Link display, provide the following information:

```
NWSSTG (Name)
NWSD (Name)
DYNAMIC (*YES)
DRVSEQNBR (*CALC)
```

Deleting network server descriptions for an IBM i logical partition that uses i virtual I/O resources

You can delete the IBM i network-server description (NWSD) for an i logical partition that uses i virtual I/O resources. When you delete the NWSD, all the configuration information for the i logical partition that uses i virtual I/O resources is deleted.

Before you begin

Before you start, ensure that you remove the disk from the auxiliary storage pool (ASP). For instructions, see [Removing a disk unit from an auxiliary storage pool \(ASP\)](#).

About this task

To delete the network-server description (NWSD) for an IBM i logical partition that uses i virtual I/O resources, follow these steps:

Procedure

1. On an IBM i control language (CL) command line on the i logical partition that provides virtual I/O resources, type the command WRKNWSD and press Enter.
2. Type 8 in the Opt field to the left of the Network Server and press Enter.
3. In the Work with Configuration Status display, if the status of the NWSD is not varied off, type 2 in the Opt field to the left of the Network Server and press Enter. Otherwise, go to the next step.
4. Press F3 to return to the previous display
5. Enter a 4 in the Opt field to the left of the Network Server and press Enter.
6. On the Confirm Delete of Network Server Descriptions display, press Enter.

Deleting virtual disk drives for an IBM i logical partition that uses i virtual I/O resources

You can delete a virtual disk drive from an IBM i logical partition that uses i virtual I/O resources. Deleting the virtual disk drive makes the space available once more to the i logical partition that provides the virtual disk resources. When you delete a virtual disk drive, all of the information on the virtual disk drive is erased.

Before you begin

Before you can delete a virtual disk drive, you must unlink the virtual disk drive from the network-server description (NWSD). For instructions, see [“Unlinking virtual disk drives from an IBM i logical partition that uses i resources” on page 192](#).

About this task

To delete a virtual disk drive, follow these steps:

Procedure

Delete the disk drive using the interface that you prefer.

Interface	Actions
IBM Navigator for i	- Click Network > Windows Administration > Disk Drives. - Right-click the disk drive that you want to delete. - Click Delete in the confirmation window.
IBM i character-based interface	- At an IBM i control language (CL) command line, type DLTNWSSTG and press F4. - Type the name of the disk drive in the Network-server storage space field and press Enter.

Using IPL types when running an IBM i logical partition that uses IBM i virtual I/O resources

The IPL source (IPLSRC) parameter on the network-server description (NWSD) determines the initial program that is loaded when the NWSD is varied on. For an IBM i logical partition that uses IBM i virtual I/O resources, the initial program is the load source. Ensure that the IPLSRC parameter specifies the location of the load source for the IBM i logical partition that uses IBM i virtual I/O resources.

You can set the IPLSRC parameter when you use the Create Network Server Description (CRTNWSD) command, and you can change the IPLSRC parameter when you use the Change Network Server Description (CHGNWSD) command.

The IPLSRC parameter has the following valid values.

IPLSRC values	Description
*Panel	The logical partition is started from the source indicated on the control panel.
*NWSSTG (network-server storage space)	This IPL type is used to start a logical partition from a virtual disk. The open firmware will find the load source in the virtual disk. If the firmware does not find the load source in the virtual disk, then the logical partition IPL will fail.
*STMF (stream file)	This IPL type is used to start a logical partition from a stream file loaded in the IBM i integrated file system of the IBM i logical partition that provides virtual I/O resources. Note that the integrated file system includes files on the optical (CD) drive on the IBM i logical partition that provides virtual I/O resources.

Unlinking virtual disk drives from an IBM i logical partition that uses i resources

By unlinking virtual disk drives (network-server storage spaces) from an IBM i logical partition that uses i resources, you disconnect the virtual disk drives from the logical partition, making the virtual disk drives inaccessible to users. If you delete an i logical partition that uses i resources, you must unlink all virtual disk drives from the logical partition before you delete the logical partition.

About this task

To unlink a virtual disk drive from an IBM i logical partition that uses i resources, follow these steps:

Procedure

Unlink disk drives from a logical partition using the interface that you prefer.

Interface	Actions
IBM Navigator for i	- Vary off the NWSD for your logical partition. - Click Network > Windows Administration > Disk Drives. - Right-click the name of the disk drive that you want to unlink. - Click Remove Link. - Select a server from the list of linked servers. - If you are unlinking a disk drive that you plan to relink later, clear Compress link sequence. You must relink the disk drive as the same link sequence number before you vary on the server. By preventing compression of the link sequence values, you avoid having to unlink and relink all the disk drives to get them in the correct sequence. - Click Remove.
IBM i character-based interface	- Vary off the NWSD for your logical partition. - At an i command line on the i logical partition that provides virtual I/O resources, type RMVNWSTGL and press F4. - In the Network-server storage space field, type the name of the storage space that you want to unlink and press Enter. - In the Network server description field, type the name of the server from which you want to unlink the storage space and press Enter. - If you are unlinking a linked disk drive that you plan to relink later, specify *NO in the Renumber field. Note: You must relink the disk drive as the same sequence number before you vary on the server. By preventing automatic renumbering, you avoid having to unlink and relink all the disk drives to get them in the correct sequence. - Press Enter. Note: If you are uninstalling a logical partition, your next step is to delete the disk drive. For instructions, see “Deleting virtual disk drives for an IBM i logical partition that uses i virtual I/O resources” on page 191. Otherwise, vary on the NWSD for your logical partition.

Saving IBM i server objects in i

When an IBM i logical partition uses i resources, the i logical partition that provides the virtual I/O resources stores the information for the i logical partition that uses the virtual I/O resources in i objects. i can restore the objects correctly only if you save all objects for an i logical partition that uses i virtual I/O resources.

You can save these objects by using options of the i GO SAVE command in the server.

- Option 21 saves the entire server.
- Option 22 saves system data, which includes objects in the QUSRSYS library.
- Option 23 saves all user data, which includes objects in the QFPNWSSTG library.

You can also use any save or restore command, or any Backup, Recovery, and Media Services (BRMS) function to save i server objects.

If you want to save a particular object, use the following table to see the location of that object on i and the command to use.

Table 29. Objects to save for logical partitions with virtual disk				
Object content	Object name	Object location	Object type	Save command
Guest partition and virtual disk drive	stgspc	/QFPNWSSTG	User-defined network-server storage spaces in system auxiliary storage pool (ASP)	GO SAV, option 21 or 23
				SAV OBJ('/QFPNWSSTG/stgspc') DEV('/QSYS.LIB/TAP01.DEVD')
			User-defined network-server storage spaces in user ASP	SAV OBJ('/QFPNWSSTG/stgspc') ('/dev/QASPnn /stgspc.UDFS')) DEV('/QSYS.LIB/TAP01.DEVD')

Table 30. Objects to save for all logical partitions with a server				
Object content	Object name	Object location	Object type	Save command
Messages from the logical partition	Various	Various	Server message queue	GO SAVE, option 21 or 23
				SAVOBJ OBJ(msg) LIB(qlibrary) DEV(TAP01) OBJTYPE(*MSGQ)
IBM i configuration objects for logical partitions	Various	QSYS	Device configuration objects	GO SAVE, option 21, 22, or 23
				SAVOBJ DEV (TAP01)
Various	Various	QUSRSYS	Various	GO SAVE, option 21 or 23
				SAVLIB LIB(*NONSYS) or LIB(*ALLUSR)

You can save and restore network server storage spaces on the i logical partition that uses i virtual resources; however, you cannot save and restore individual files.

Backing up and recovering IBM i logical partitions that use i virtual I/O resources

You can back up and recover an IBM i logical partition that uses resources from another i logical partition by using the GO SAVE operation.

Managing Linux logical partitions that use IBM i resources

You can manage Linux logical partitions that uses IBM i virtual I/O resources to help maximize utilization of the physical hardware and simplify the backup procedure for your managed system.

Adding virtual disk units to a Linux logical partition

You can add virtual disk units dynamically to a Linux logical partition that uses IBM i resources. This allows you to increase the storage capacity of your AIX logical partition when needed.

About this task

Virtual disks simplify hardware configuration on the server because they do not require you to add additional physical devices to the server in order to run Linux. You can allocate up to 64 virtual disks to a Linux logical partition. Each virtual disk supports up to 1000 GB of storage. Each virtual disk appears to Linux as one actual disk unit. However, the associated space in the i integrated file system is distributed across the disks that belong to the i logical partition. Distributing storage across the disks provides the benefits of device parity protection through i. Therefore, you do not have to use additional processing resources and memory resources by setting up device parity protection through Linux.

IBM i provides the ability to dynamically add virtual disks to a Linux logical partition. You can allocate disk space in the integrated file system and make it available to Linux without restarting the server or logical partition. The Linux administrator can also configure the newly allocated disk space and make it available without restarting the server.

To add virtual disks dynamically to a Linux logical partition, complete the following steps:

Procedure

1. If you use IBM Navigator for i, create a network-server storage space using IBM Navigator for i.
 - a) Expand **My Connections > your server > Network > Windows Administration**.
 - b) Right-click the **Disk Drives** and select **New Disk**.
 - c) In the **Disk drive name** field, specify the name that you want to give to the network-server storage space.
 - d) In the **Description** field, specify a meaningful description for the network-server storage space.
 - e) In the **Capacity** field, specify the size of the new network-server storage space in megabytes.
Refer to the installation documentation of your preferred Linux distributor to determine the size you want to use.
 - f) Click **OK**.
 - g) Continue with step “4” on page 195.
2. If you use a character-based interface, create a network-server storage space using the character-based interface:
 - a) At an IBM i command line, type the command CRTNWSSTG and press F4.
The Create NWS Storage Space (CRTNWSSTG) display opens.
 - b) In the Network-server storage space field, specify the name you want to give to the network-server storage space.
 - c) In the Size field, specify the size in megabytes for the new network-server storage space.
Refer to the installation documentation of your preferred Linux distributor to determine the size you want to use.
 - d) In the Text description field, specify a meaningful description for the network-server storage space.
 - e) Press Enter.
3. If you use IBM Navigator for i, link the network-server storage space using IBM Navigator for i.
 - a) Expand **My Connections > your server > Network > Windows Administration**.
 - b) Click **Disk Drives**, right-click an available network-server storage space, and select **Add Link**.
 - c) Select the server to which you want to link the network-server storage space.
 - d) Select one of the available data access types.
 - e) Click **OK**.
 - f) Continue with step “5” on page 196.
4. If you use a character-based interface, link the network-server storage space using a character-based interface:
 - a) At an IBM i command line, type the command ADDNWSSTGL and press F4.
The Add Network-Server Storage Link (ADDNWSSTGL) display opens.
 - b) In the Network server description field, specify the name of the network server description (NWSD).
 - c) In the Dynamic storage link field, specify *YES to make the network-server storage space dynamically available to the logical partition (that is, available without rebooting the Linux logical partition).
 - d) In the Drive sequence number field, specify the link sequence position you want to use.

- e) Press Enter.
- If the Linux logical partition is not running, activate the Linux logical partition. Do not continue until the logical partition is running.
 - Log in to Linux using a user name with superuser (root) privileges.
 - Determine the host ID, SCSI bus, and logical unit number (LUN) for your new virtual disk drive.

You can list the existing devices by typing the following command at the Linux command prompt:
`cat /proc/scsi/scsi`. The following example shows sample output of the command:

```
Attached devices:
Host: scsi0 Channel: 00 Id: 00 Lun: 00
Vendor: IBM      Model: VDASD NETSPACE   Rev: 0001
Type:   Direct-Access          ANSI SCSI revision: 04
```

In this example, NETSPACE is the name of the network storage space for the displayed device. Look for the name of an existing network storage space on your Linux logical partition. Note the numeric part of the `Host:` value (host ID) and the `Channel:` (SCSI bus) and `Lun:` (logical unit number (LUN)) values for the existing network storage space. The new virtual disk drive will have the same host ID, SCSI bus, and LUN as the existing network storage space. For example, if the existing network storage space is as displayed in the preceding example output, then the new virtual disk drive will have a host ID of 0, a SCSI bus of 0, and a LUN of 0.

- Determine the SCSI ID for your new virtual disk drive.

You can list the existing devices in table form by typing the following commands at the Linux command prompt:

```
cd /proc/scsi/sg
cat device_hdr; cat devices
```

The following example shows sample output of the commands:

host	chan	id	lun	type	opens	qdepth	busy	online
0	0	0	0	0	2	30	0	1
0	1	0	0	0	0	30	0	1

Note the `host` (host ID), `chan` (SCSI bus), `id` (SCSI ID), and `lun` (logical unit number (LUN)) values for the existing devices. Find the devices that have the same host ID, SCSI bus, and LUN as the new virtual disk drive (as you determined in the previous step). Of those devices, find the device with the greatest SCSI ID. The new virtual disk drive will have a SCSI ID that is one greater than the greatest existing SCSI ID. For example, if the new virtual disk drive has a host ID of 0, a SCSI bus of 0, and a LUN of 0, and the devices on your Linux logical partition are as listed in the example output above, then the new virtual disk drive will have a SCSI ID of 1.

- Type the following command at the Linux command prompt to add the virtual disk drive manually:
`echo "scsi add-single-device host chan id lun" > /proc/scsi/scsi`.

Use the following information to help you understand the arguments of the command:

- `host` is the host ID.
- `chan` is the SCSI bus.
- `id` is the SCSI ID.
- `lun` is the LUN.

For example, if the new virtual disk drive is to have a host ID of 0, a SCSI bus of 0, a SCSI ID of 1, and a LUN of 0, you would type the command `echo "scsi add-single-device 0 0 1 0" > /proc/scsi/scsi` at the Linux command prompt.

- At the Linux command prompt, type the following command to create a disk partition on the virtual disk drive: `fdisk /dev/sdb`.

You must have superuser (root) privileges to run this command.

The Command (`m` for `help`): prompt is displayed.

- Type `p` at the prompt to see the current partition table for the virtual disk drive.

By default, the new virtual disk drive shows a single disk partition on the virtual disk.

For example,

```
Disk /dev/sdb: 64 heads, 32 sectors, 200 cylinders
Units = cylinders of 2048 * 512 bytes

Device Boot      Start         End      Blocks   Id  System
/dev/sdb1          1          199       203760    6   FAT16
```

12. Type `d` at the command prompt to delete the current partition and then create a new one.

The default format for the disk partition is FAT16. Do not use a disk partition that is formatted as FAT16 on your virtual disk drive.

The Partition number (1-4): prompt is displayed.

13. Type the disk partition number you want to delete and press Enter.

In this example, you type a 1.

The `fdisk` command indicates that the deletion is successful by displaying the command prompt.

14. Type `n` to create a new disk partition.

The Command action E extended P primary partition (1-4) prompt is displayed.

15. Type `p` to create a primary disk partition on the virtual disk and press Enter.

The Partition number (1-4): prompt is displayed.

16. Type `1` because this is the first partition on the virtual disk, and press Enter.

The First cylinder (1-200, default 1): prompt is displayed.

17. Press Enter to use the default of 1 for the first disk cylinder.

This uses the entire disk for this disk partition.

The Last cylinder or +size or +sizeM or +sizeK (1-200, default 200): prompt is displayed.

18. Press Enter to use the default of 200 for the last disk cylinder.

This uses the entire virtual disk for this partition.

Note: The type of the partition defaults to Linux. If you need a different disk type (like Logical Volume Manager (LVM), or Linux Extended), type `t` to change the type of the partition.

The `fdisk` command indicates that the partition creation is successful by returning the command prompt.

19. Type `w` to commit the changes to the disk structure and press Enter.

The `fdisk` command writes the changes to the virtual disk drive. The `fdisk` command displays the following diagnostic message:

```
The partition table has been altered!

Calling ioctl() to re-read partition table.
Syncing disks.
```

After the operation is completed, the `fdisk` command returns the command prompt.

20. Format the disk partition using the Linux **mkfs** command .

There are a number of optional parameters for the `mkfs` command, but typically the defaults satisfy most disk uses. To format the disk partition created in the previous steps, ensure that you are logged in with superuser (root) privileges and type the following command at a Linux command prompt:

```
mkfs /dev/sdb1
```

Since a single disk partition exists on the second virtual disk, the name of the disk is `/dev/sdb1` (the `sdb` indicates that it is the second disk, and the `1` indicates that it is partition 1). The `mkfs` command displays the following diagnostic messages:

```
mke2fs 1.28 (31-Aug-2002)
Filesystem label=
OS type: Linux Block size=1024 (log=0)
```

```

Fragment size=1024 (log=0)
51200 inodes, 204784 blocks
10239 blocks (5.00%) reserved for the super user
First data block=1
25 block groups
8192 blocks per group, 8192 fragments per group
2048 inodes per group
Superblock backups stored on blocks:
    8193, 24577, 40961, 57345, 73729

Writing inode tables: done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 29 mounts or
180 days, whichever comes first. Use tune2fs -c or -i to override.

```

21. Type the following command to create a directory that you can use to access the new file:
`mkdir /mnt/data`
22. Type the following command to mount the virtual disk drive in the new directory: `mount /dev/sdb1 /mnt/data`
23. Add an entry to the `/etc/fstab` file using a Linux text editor, such as `vi`.
For example, `/dev/sdb1 /mnt/data ext2 defaults 1 1`. This entry mounts the virtual disk every time you restart Linux.

Linking a network-server storage space to a network-server description

You can link a network-server storage space (NWSSTG) to one or more network-server descriptions (NWSDs). This allows the NWSDs and their associated logical partitions to use the data stored on the NWSSTG.

About this task

You can link an NWSSTG to an unlimited number of NWSDs. This is beneficial when multiple logical partitions need access to a single application.

When you link an NWSSTG to an NWSD, you can set up the NWSD to have read-only access to the NWSSTG, or you can set up the NWSD to read or write to the NWSSTG.



Attention: If more than one NWSD can write to the NWSSTG, ensure that only one NWSD can update the data at a time. Otherwise, changes made by one NWSD can be overwritten by another NWSD.

To link an NWSSTG to an NWSD, follow these steps:

Procedure

1. At an IBM i command line, type the command `ADDNWSSTGL` and press F4.
2. From the Add Server Storage Link display, provide the following information:

```

NWSSTG (Name)
NWSD (Name)
DYNAMIC (*YES)
DRVSEQNBR (*CALC)

```

3. Press F10 (Additional Parameters).
4. Enter the type of access the storage space will have.

Deleting network server descriptions for a Linux logical partition

You can delete the IBM i network server description (NWSD) for a Linux logical partition that uses i resources. When you delete the NWSD, all the configuration information for the Linux logical partition is deleted from i.

About this task

To delete the network-server description (NWSD) for a Linux logical partition, follow these steps:

Procedure

1. On an i control language (CL) command line, type the command WRKNWSD and press Enter.
2. Type 8 in the Opt field to the left of the Network Server and press Enter.
3. In the Work with Configuration Status display, if the status of the NWSD is not varied off, type 2 in the Opt field to the left of the Network Server and press Enter. Otherwise, go to the next step.
4. Press F3 to return to the previous display
5. Enter a 4 in the Opt field to the left of the Network Server and press Enter.
6. On the Confirm Delete of Network Server Descriptions display, press Enter.

Deleting virtual disk drives for a Linux logical partition

You can delete a virtual disk drive from a Linux logical partition that uses IBM i resources to make the space available to the i logical partition once more. When you delete a virtual disk drive, all of the information on the virtual disk drive is erased.

Before you begin

Before you can delete a disk drive, you must unlink it from the network-server description. For instructions, see [“Unlinking virtual disk drives from a Linux logical partition” on page 200.](#)

About this task

To delete a virtual disk drive, follow these steps:

Procedure

Delete the disk drive using the interface that you prefer.

Interface	Actions
IBM Navigator for i	Complete the following steps: a. Click Network > Windows Administration > Disk Drives. b. Right-click the disk drive that you want to delete. c. Click Delete. d. Click Delete in the confirmation window.
IBM i character-based interface	Complete the following steps: a. At an IBM i control language (CL) command line, type DLTNWSSTG and press F4. b. Type the name of the disk drive in the Network-server storage space field and press Enter.

Using IPL types when running Linux

The IPL source (IPLSRC) parameter on the network server description (NWSD) determines the initial program that is loaded when the NWSD is varied on. For a Linux logical partition that uses IBM i resources, the initial program is the kernel. Ensure that the IPLSRC parameter specifies the kernel location of the kernel for the Linux logical partition that uses i resources.

You can set the IPLSRC parameter when you use the Create Network Server Description (CRTNWSD) command, and you can change the IPLSRC parameter when you use the Change Network Server Description (CHGNWSD) command.

Note: The IPLSRC parameter also has the values A, B, and D, which are not valid for hardware that is used by IBM i logical partitions.

The IPLSRC parameter has the following valid values.

IPLSRC values	Description
*Panel	The logical partition is started from the source indicated on the control panel.
*NWSSTG (network-server storage space)	This IPL type is used to start a logical partition from a virtual disk. The open firmware will find the kernel in the virtual disk. The open firmware searches the first virtual disk connected to the server for a logical partition marked bootable, and of type 0x41 (PReP start). If a logical partition of this type does not exist, the logical partition IPL will fail.
*STMF (stream file)	This IPL type is used to start a logical partition from a kernel loaded in the IBM i integrated file system. Note that the integrated file system includes files on the optical (CD) drive on IBM i.

Unlinking virtual disk drives from a Linux logical partition

By unlinking virtual disk drives (network-server storage spaces) from a Linux logical partition that uses IBM i resources, you disconnect the virtual disk drives from the logical partition, making the virtual disk drives inaccessible to users. If you delete a Linux logical partition that uses i resources, you must unlink all virtual disk drives from the logical partition before you delete the logical partition.

About this task

To unlink a virtual disk drive from a Linux logical partition that uses i resources, follow these steps:

Procedure

- Unlink disk drives from a logical partition using IBM Navigator for i.

If you prefer to use a character-based interface, go to step “2” on [page 200](#).

 - Vary off the NWSD for your logical partition.
 - Click **Network > Windows Administration > Disk Drives**.
 - Right-click the name of the disk drive that you want to unlink.
 - Click **Remove Link**.
 - Select a server from the list of linked servers.
 - If you are unlinking a disk drive that you plan to relink later, uncheck **Compress link sequence**.
You must relink the disk drive as the same link sequence number before you vary on the server. By preventing compression of the link sequence values, you avoid having to unlink and relink all the disk drives to get them in the correct sequence.
 - Click **Remove**.
 - You have completed this procedure. Do not complete step “2” on [page 200](#).
- Unlink disk drives from a logical partition using a character-based interface:
 - Vary off the NWSD for your logical partition.
 - Type RMVNWSSTGL and press F4.
 - In the Network-server storage space field, type the name of the storage space that you want to unlink and press Enter.
 - In the Network server description field, type the name of the server from which you want to unlink the storage space and press Enter.
 - If you are unlinking a linked disk drive that you plan to relink later, specify *NO in the Renumber field.

Note: You must relink the disk drive as the same sequence number before you vary on the server. By preventing automatic renumbering, you avoid having to unlink and relink all the disk drives to get them in the correct sequence.

f) Press Enter.

Note: If you are uninstalling a logical partition, your next step is to delete the disk drive. For instructions, [“Deleting virtual disk drives for a Linux logical partition” on page 199](#). Otherwise, vary on the NWSD for your logical partition.

Saving Linux server objects in IBM i

When a Linux logical partition uses IBM i resources, i stores Linux information in i objects. IBM i can restore the objects correctly only if you save all objects for a Linux logical partition.

You can save these objects by using options of the i GO SAVE command in the server.

- Option 21 saves the entire server.
- Option 22 saves server data, which includes objects in the QUSRSYS library.
- Option 23 saves all user data, which includes objects in the QFPNWSSTG library.

If you want to save a particular object, use the following table to see the location of that object on i and the command to use.

<i>Table 31. Objects to save for logical partitions with virtual disk</i>				
Object content	Object name	Object location	Object type	Save command
Guest partition and virtual disk drive	stgspc	/QFPNWSSTG	User-defined network-server storage spaces in system auxiliary storage pool (ASP)	GO SAV, option 21 or 23
				SAV OBJ('/QFPNWSSTG/stgspc') DEV('/QSYS.LIB/TAP01.DEVD')
			User-defined network-server storage spaces in user ASP	SAV OBJ('/QFPNWSSTG/stgspc') ('/dev/QASPnn /stgspc.UDFS') DEV('/QSYS.LIB/TAP01.DEVD')

<i>Table 32. Objects to save for all logical partitions with a server</i>				
Object content	Object name	Object location	Object type	Save command
Messages from the logical partition	Various	Various	Server message queue	GO SAVE, option 21 or 23
				SAVOBJ OBJ(msg) LIB(qlibrary) DEV(TAP01) OBJTYPE(*MSGQ)
IBM i configuration objects for logical partitions	Various	QSYS	Device configuration objects	GO SAVE, option 21, 22, or 23
				SAVOBJ DEV (TAP01)
Various	Various	QUSRSYS	Various	GO SAVE, option 21 or 23
				SAVLIB LIB(*NONSYS) or LIB(*ALLUSR)

Backing up and recovering Linux logical partitions that use IBM i virtual I/O resources

When you create a Linux logical partition that uses resources from an IBM i logical partition, you can manage backup and recovery using IBM i control language (CL) commands, Linux commands, or a combination of the two.

To save Linux data in a logical partition that uses i resources to a shared tape drive and restore the data from the tape drive, you can use either the Linux **tar** command or the i Save (SAV) and Restore (RST) commands. You can also use the **tar** command to save your data to a file. If you use the **tar** command to save data, the only way you can restore that data is by using the **tar** command again. Similarly, if you use the SAV command to save data, the only way you can restore that data is by using the RST command. The two methods of backing up and restoring data are not compatible.

The following restrictions apply:

- To use the tape device from Linux, you must vary the tape off under IBM i.
- Saving the storage space is typically faster than saving by using the **tar** command, but it does not provide file-level backup and recovery.
- Linux does not support switching tapes in a library device. You can only use the tape that is currently in the device.
- You cannot save IBM i data and **tar** data on the same tape volume.

Backing up and recovering files using the tar command

The most common data backup utility in Linux is the **tar** (tape archive) utility. Use the Linux **tar** command if you have Linux installed on a dedicated disk or if you cannot vary off a Linux logical partition while you are backing up data.

Backups using the Linux **tar** command are at the file level. They save only the files and directories that the **tar** command specifies. Therefore, you cannot use the **tar** command to save Linux data that is not in the file server. For example, you cannot save a kernel in the PowerPC® Reference Platform (PReP) start logical partition by using the **tar** command.

One advantage of the **tar** command is that it supports incremental backups and backup of special devices, which is not common for **tar** implementations. Also, the **tar** command backs up files without regard to the underlying file system type.

Saving to and restoring from a tape device

Use these procedures to save and restore Linux files between a Linux logical partition that uses IBM i resources and a shared tape drive.

Before you begin

Ensure that your Linux data is in the file server.

Linux typically treats tape as a *character device* that it can quickly read from or write to in long streams of data, but cannot quickly access to find specific data. By contrast, Linux treats a disk or CD as a *block device* that it can read from or write to quickly at any point on the device, making it suitable for the **mount** command.

About this task

Complete the following steps to save and restore Linux files between a logical partition that uses IBM i resources and a shared tape drive:

Procedure

1. Type the following command: `tar -b 40 -c -f /dev/st0 files`

Use the following descriptions to help you understand the arguments of this command:

- **tar** is the command name (the contraction of "tape archive").

- **-b 40** is the block size in sectors. This argument specifies that Linux is to write the archive stream in blocks of 40 sectors (20 KB). If you do not specify a value for this argument, the default value is 20 sectors (10 KB), which does not perform as well over virtual tape as does a value of 40.
- **-c** is the command action to create. This argument specifies that the **tar** command creates a new archive or overwrites an old one (as opposed to restoring files from an archive or adding individual files to an existing archive).
- **-f /dev/st0** is the virtual tape device and number. This argument specifies that the command uses virtual tape 0 on the server. After the **tar** command runs, the tape device is closed and the tape is rewound. To save more than one archive on the tape, you must keep the tape from rewinding after each use, and you must position the tape to the next file marker. To do this, specify the *nst0* (nonrewinding virtual tape) device instead of *st0*.
- *files* are the names of the files and directories that you plan to save.

You have now saved Linux data from a logical partition that uses IBM i resources to the shared tape drive.

2. Type the following command: `tar -b 40 -x -f /dev/st0 files`

The **-x** (extract) argument replaces the **-c** (create) argument in the **tar** command used in step “1” on page 202.

You have now restored Linux data from the shared tape drive to a logical partition that is sharing resources.

Saving to and restoring from a file

You can save and restore Linux files between a Linux logical partition that uses IBM i resources and a tar file.

Saving to a file

The following is an example of using the **tar** command to save to a file.

```
tar -cvf /tmp/etc.tar /etc
```

Use the following descriptions to help you understand the arguments of this command:

tar

The command name.

c

Create a tar file.

v

Verbose. This argument shows the files that are being added to the tar file.

f

The data immediately following **f** is the name of the tar file.

/tmp/etc.tar

The name of the tar file.

/etc

An object to be added to the tar file. Because */etc* is a directory, the utility adds all the contents of the directory and its subdirectories to the tar file.

After you create the tar file, you can save it to an offline medium in several ways. For example, you can save the tar file to a virtual tape device or a directly attached tape device. You can also copy the tar file to the integrated file system and save it at a later time.

You can save the data on a Linux logical partition to a tar file during normal server usage. You can automate and start the **tar** utility by using the **cron** (chronology) daemon on the logical partition. The **cron** daemon is a scheduling mechanism for Linux. You can also use the **tar** utility to schedule a single backup request. For example, if you want to use the tar utility to back up the */etc* directory at 10 p.m. on 19 September, you can type the following command: `at 10pm Sep 19 -f tar.command.`

Restoring from a file

The following is an example of using the **tar** command to restore from file: `tar -xvf /tmp/etc.tar /etc`. The **-x** (extract) argument replaces the **-c** (create) argument in the **tar** command used to save the files.

Backing up and recovering Linux logical partitions using i commands

If you have a Linux logical partition that uses IBM i resources, tools are available in i for backup and recovery. You can use the Save (SAV) and Restore (RST) control language (CL) commands to save and restore entire virtual disks in their current state.

The SAV command saves the directory that has the same name as the virtual disk under the QFPNWSSTG directory in the integrated file system. This method of backup and recovery is most effective if the Linux kernel is saved in a PowerPC Reference Platform (PReP) start logical partition on the virtual disk. On most Linux distributions, this usually occurs as part of a default installation.

Backups of storage spaces using i commands are at drive level. This means that i backs up the entire contents of a virtual disk, or network storage space, rather than individual files. Thus, the correct SAV command backs up any information on the drive, including a kernel in the PReP start logical partition.

If you save the Linux kernel in a PReP logical partition, you can restore and start the logical partition after a total system re installation. You can also transport and restore saved virtual disks to other servers using File Transfer Protocol (FTP) and tape.

Save Linux data by using IBM i SAV

You can save data for a Linux logical partition that uses IBM i resources by using the Save (SAV) i CL command.

About this task

On IBM i, your data is in a network-server storage space.

Saving and restoring individual Linux files using IBM i commands requires that you use the QNTC directory in the integrated file system. You can access the files that you save and restore by using a file share. You can define the file share by using Samba on Linux, and you can access the file share by using QNTC.

To save data for a Linux logical partition that uses IBM i resources by using the Save (SAV) IBM i CL command, follow these steps:

Procedure

1. At the IBM i command line, enter the Save (SAV) command.
2. On the Save display, enter the following parameter values:
 - a) In the **Device** field, enter the associated IBM i device description.
To save to a file in a library like QGPL, enter `/qsys.lib/qgpl.lib/myfile.file`. For example, if your tape device is named TAP01, enter `/qsys.lib/tap01.devd`.
 - b) In the **Objects: Name** field, enter the server, share, or file.
For example, if your server is named MYSERVER, your share is named MYSHARE, and it contains all of the directories and files that need to be saved, enter `/QNTC/MYSERVER/MYSHARE`.
3. At the IBM i command line, enter the Display Save File (DSPSAVF) command to verify that the changed save file exists.
4. In the Option field by the new save file name, enter 5 (Display) to display a list of the stream files in the save file.

Restore Linux data using i RST

You can restore data for a Linux logical partition that uses i resources by using the Restore (RST) IBM i CL command.

About this task

Saving and restoring individual Linux files using IBM i commands requires that you use the QNTC directory in the integrated file system. You can access the files that you save and restore by using a file share. You can define the file share by using Samba on Linux, and you can access the file share by using QNTC.

Restore (RST) is the i command to restore Linux files from the shared tape drive of the logical partition that shares resources. On the Restore Object display, enter the following parameter values:

Procedure

1. To restore from a tape device, enter the associated i device description in the **Device** field.
For example, if your tape device is named TAP01, enter `/qsys.lib/tap01.devd`.
2. To restore from a save file in library QGPL, enter the associated file name.
For example, `/qsys.lib/qgpl.lib/myfile.file`.
3. In the **Objects: Name** field, enter the server, share, or file.
For example, if your server is named MYSERVER, your share is named MYSHARE, and it contains all of the directories and files that need to be restored, enter `/QNTC/MYSERVER/MYSHARE`.

Backing up the network server description and virtual disk drives associated with a Linux logical partition
Learn about how to back up the data for a Linux logical partition that uses IBM i resources.

Backing up the data for a Linux logical partition that uses IBM i resources is different from backing up the data for a Linux logical partition that uses its own resources. When you install the logical partitions with virtual disk, the IBM i logical partition that shares resources creates a network server description and creates disk drives for your Linux logical partition that you need to back up. Some of the disk drives are server-related (the installation and server drives), while others are user-related. Because your Linux logical partition might consider the disk drives to be a unified server, you must save all the disk drives and the network server description so they restore correctly.

With the implementation of a logical partition, you can save and restore virtual disks as IBM i network-server storage space objects. These objects are saved as part of the server when you perform a full server backup. You can also specifically save the network server description and storage spaces that are associated with a logical partition on a server. Daily backup of the server drive is a good practice.

Building a rescue image on a network storage space

You can build a rescue image on a network storage space (NWSSTG) to assist you in checking and repairing a faulty Linux installation.

Before you begin

A *rescue image* is a disk image that contains the Linux kernel, a shell, and the diagnostic tools, drivers, and other utilities that would be useful for checking and repairing a faulty Linux installation. Many Linux distributors include a rescue image on their installation disks. One rescue solution for a logical partition is to create a small NWSSTG that can remain on the integrated file system solely for the purpose of rescuing logical partitions. You can install a rescue image to the NWSSTG when you create your logical partition.

Before creating a rescue image on network storage, it is important to document the configuration information for each of your logical partitions.

1. Document the drive configuration information, which is located in the `/etc/fstab` file.
2. Capture the networking information that is reported when you run the **ifconfig** command.
3. Create a list of the modules that are needed by each logical partition. You can see which modules are in use by using the **lsmod** command from within Linux. Use the information obtained from the

commands and files listed above to determine which files to store on your rescue network storage space.

About this task

To build a rescue image on an NWSSTG, follow these steps:

Procedure

1. Determine how much network storage space you need to build the rescue image.
Consult your Linux documentation to see how much space is required for a minimum installation of your distribution, and add enough space to create a swap partition (a PowerPC Reference Platform (PReP) start partition) and to install any extra software that you would like to have available in your rescue image. For example, if the documentation states that a minimum server installation is 291 MB, create a storage space of 425 MB.
2. Create a network storage space (CRTNWSSTG) of the size you determined for the rescue image.
You might want to make a note in the storage space description field that indicates which distribution was used to make the rescue image and warns that it should be saved.
3. Link this storage space to a network server description (NWSD).
You do not need to create a new NWSD for this step. You could unlink an existing storage space and temporarily link your rescue storage space to any of your existing NWSDs.
4. Start the installation server for your distribution as described in the documentation and follow the prompts.
To partition your installation manually, ensure that you create a PReP start partition. At the point where you select the packages to install, select the minimum number of packages supported. The name for the package group varies by distribution.
5. Allow the installer to complete its package installation and configuration.
After installation has finished, the installer starts the rescue image for you.
6. Verify that the rescue image has all the utilities that you need.
For a logical partition, at a Linux command prompt, type `rpm -qa | grep ibmsis` to make sure that the utilities that work with the integrated disk are available.
7. Ensure that the device drivers that your logical partitions require are installed.
For example, verify that `pcnet32` is installed for Ethernet devices, or that `olympic` is installed for token-ring devices. The kernel modules that have been compiled can be found in the `/lib/modules/kernel version/kernel/drivers` directory or in directories under that directory.
8. Install any other special drivers or software packages that your logical partitions require.
9. Use File Transfer Protocol (FTP) to send the files with the configuration information for your other logical partitions to the rescue server network storage space.
10. Install the kernel manually (if you are required to do so by your Linux distribution).
For details regarding installing the kernel, consult the appropriate installation documentation for your distribution.
11. Make note of the path to the root partition on the rescue-storage space.
You must use this information to start the rescue network storage space from the network. To determine the root partition, type the command `cat /etc/fstab`. The partition that has a forward slash (/) in the second column is your root partition. For further assistance in determining the root partition, see the documentation for your distribution.

What to do next

You can shut down your logical partition by typing `shutdown -h now` and varying off the logical partition after the shutdown has completed. After the logical partition has varied off, you can unlink the rescue storage space and relink the normal storage space for the NWSD.

Using a rescue image from a network-server storage space

You can use a Linux rescue image on a network-server storage space (NWSSTG) to repair a Linux logical partition that uses IBM i resources. A *rescue image* is a disk image that contains the Linux kernel, a shell, and the diagnostic tools, drivers, and other utilities that would be useful for checking and repairing a faulty Linux installation.

About this task

To use the rescue image that you built on the NWSSTG, use the following steps:

Procedure

1. Disconnect the virtual storage space for the failed logical partition (if applicable) by using the Work with NWS Storage Spaces (WRKNWSSTG) command.
2. Connect your rescue storage space as the first drive to the network server description (NWSD), and reconnect the original storage space (where applicable) as the second drive.
3. Edit the NWSD for the failed partition so that it starts from IPL source *NWSSTG. Also, edit the IPL Parameters field to reflect the root partition on the rescue storage space. For most distributions, this is a parameter such as `root=/dev/sda3` or `root=/dev/vda1`. For assistance, see the documentation for your Linux distribution.
4. Restart the partition.
5. If the existing root partition is on a dedicated disk, you might need to insert the `ibmisis` driver using the `insmod ibmisis` command.
6. Create a mount point to which you will mount the root partition of the network storage space that you are trying to rescue. You can use a command such as `mkdir /mnt/rescue`.
7. Mount the root partition of the network storage space that you are trying to rescue. Mount a drive using the command `mount -t partition-type partition-location mount-point`, where the partition type is the format of the partition such as `ext2` or `reiserfs`, the partition location is similar to `/dev/sdb3` (for non-devfs disk partitions), `/dev/sd/disc1/part3` (for devfs disk partitions), or `/dev/sda2` (for a partition on a dedicated disk).
8. The drive that you are trying to rescue, when using virtual disk, will be the second drive rather than the first drive. (That is, if the drive was `/dev/sda3` when the partition was running normally, it will be `/dev/sdb3` in the rescue server.)
9. Use the documentation or the configuration files you created when you created the rescue NWSSTG to help you determine the device for the root of the partition you are trying to rescue. Your mount point will be similar to `/mnt/rescue` if you use the previous example.

What to do next

You can either use the rescue tools provided in your rescue storage space against the mount point you have created or you can work on the partition that you are rescuing from within its own storage space. If rescuing the image from its own storage space, change the root directory for that partition using the `chroot mount-point` command.

Backing up network server descriptions for a Linux logical partition

When you save the storage space objects that are associated with a logical partition that uses virtual disks, you must also save the network server description (NWSD). Otherwise, a logical partition might not be able to re-establish items such as the file-system permissions for the logical partition.

About this task

Use the Save Configuration (SAVCFG) command to save the network server description:

Procedure

1. On the IBM i command line, type SAVCFG.
2. Press Enter to save the NWSD configuration.

What to do next

The Save Configuration command (SAVCFG) saves the objects associated with an NWSD, including the line descriptions and network-server storage space link information. SAVCFG does not save the storage spaces associated with this server. You can use the Save Object (SAV) command to save the storage spaces.

Restoring network-server descriptions for a Linux logical partition

In a disaster-recovery situation, you would restore all the configuration objects, which include the network-server description (NWSD) for your logical partition. In some situations, you must specifically restore the NWSD. For example, you must restore the NWSD when you migrate to new hardware.

Before you begin

To have IBM i automatically relink disk drives within the integrated file system to the restored NWSD, restore those disk drives first.

About this task

To restore the NWSD, use the Restore Configuration (RSTCFG) command:

Procedure

1. On an IBM i command line, type RSTCFG and press F4 (Prompt).
2. In the **Objects** field, specify the name of the NWSD.
3. In the **Device** field, specify which device you are using to restore the NWSD.
If you are restoring from media, specify the device name. If you are restoring from a save file, specify *SAVF and identify the name and library for the save file in the appropriate fields.
4. Press Enter to restore the NWSD.
5. When you have restored the NWSD and all of its associated storage spaces, start (vary on) the logical partition.

Synchronizing the hypervisor and Service Processor time-of-day clocks to Time Reference Partition

You can ensure that the time-of-day clocks used by the hypervisor and Service Processor are accurate through the usage of one or more Time Reference Partitions (TRP). Whenever the time-of-day changes in the TRP, the hypervisor and the service processors time-of-day are updated to match the time that is specified by the TRP. The time of the TRP can be changed manually or can be managed by using the Network Time Protocol (NTP) support. NTP can be used to automatically ensure consistent time-of-day across multiple servers. When you designate a logical partition as a TRP, you must choose a logical partition that cannot be migrated to another server such as a Virtual I/O Server (VIOS) partition. More than one TRP can be specified per server and the longest running TRP is recognized as the TRP of the system.

About this task

The service processor uses its time-of-day clock for time stamps of various error logs and events. The service processor has a battery backed clock so if an electrical outage occurs, the server can maintain the current time-of-day. Whenever a server is powered on, the hypervisor time-of-day clock is initialized from the time-of-day clock of the service processor. The hypervisor uses its time-of-day clock whenever a new partition is created. Newly created partitions start with the time-of-day clock of the hypervisor. After a logical partition is created, the time-of-day clock of the logical partition is separate from the time-of-day clock of the hypervisor (changes to the time of the hypervisor does not affect existing partitions).

To enable TRP on a partition, complete the following steps:

Procedure



1. In the navigation pane, click the **Resources** icon.
2. Click **All Partitions**. Alternatively, click **All Systems**. In the work pane, click the server name that has the logical partition. Click **View System Partitions**. The All Partitions page is displayed.
3. In the work pane, select logical partition and click **Actions > View Partition Properties**. The **Properties** page is displayed.
4. Click the **General** tab.
5. Click the **Advanced** tab. In the **Advanced Settings** area, select the **Enable Time Reference** check box.
6. Repeat steps 3 through 5 for any additional logical partitions.

Enabling user mode access to the hardware accelerator

You can enable user mode access to the hardware accelerator by using the Hardware Management Console (HMC). The HMC must be at Version 9.1.940, or later.

About this task

The GNU zip (gzip) accelerators are compression and decompression cards that are used to increase server performance and network I/O efficiency. Quality of service (QoS) credits is a mechanism that is used to give logical partitions access to shared hardware accelerators. You can verify whether the server supports enabling user mode access to the hardware accelerator by using the **lssyscfg** command. To view the supported hardware accelerators types, along with the corresponding maximum Hardware Accelerator QoS, currently available Hardware Accelerator QoS, and pending available Hardware Accelerator QoS credits of the server, you can use the **lshwres** command. You can also use the **lshwres** command to view the amount of gzip QoS that is assigned to a logical partition. You can enable the gzip quality of service (QoS) credits for an logical partition AIX, Linux, or Virtual I/O Server logical partition that are either in the Activated or Not Activated states by using the **chhwres** command. You cannot enable the QoS credits for an IBM i logical partition.

The logical partition is assigned credits only when the following conditions are met:

- The server supports enabling the Hardware Accelerator.
- The server supports the Hardware Accelerator type that you specify.
- The operating system supports enabling QoS credits dynamically.
- Sufficient amount of Hardware Accelerator QoS credits are available that can be assigned to the logical partition.

To remote restart a logical partition that is enabled with gzip QoS credits, the destination server must support user mode access to the hardware accelerator. Additionally, the following conditions apply:

- When the HMC that manages the destination server is at version 9.1.0, or earlier, the remote restart operation succeeds with the logical partition losing the gzip QoS credits after the remote restart operation completes.
- The remote restart operation succeeds when there is a sufficient amount of available gzip QoS credits on the destination server and the HMC that manages the destination server is at version 9.1.940, or later. The logical partition is restarted with the same amount of gzip QoS credits that was available at the source server.
- When the destination server has an insufficient amount of available gzip QoS credits, and the server is managed by an HMC at version 9.1.940, or later, the remote restart operation succeeds and the logical partition is assigned partial gzip QoS credits that depends on the amount of available gzip QoS credits on the destination server.

Related information

[chhwres command](#)

[lshwres command](#)

[lssyscfg command](#)

Performance considerations for logical partitions

You can manage and enhance the performance of logical partitions so that your system uses its resources in the most efficient manner.

You can manage and enhance the performance of a AIX logical partition by configuring the AIX operating system.

Managing IBM i performance ensures that your managed system is efficiently using resources and that your managed system provides the best possible services to you and to your business. Moreover, effective performance management can help you quickly respond to changes in your managed system and can save on expenses by postponing costly upgrades and service fees.

Related information

[AIX Performance management](#)

[Performance Tools Guide and Reference](#)

[Performance Toolbox Version 2 and 3 Guide and Reference](#)

[Power Systems Capacity on Demand](#)

Adjusting the Active Memory Expansion configuration to improve performance

You can run the Active Memory Expansion planning tool to generate performance statistics for an AIX logical partition that uses Active Memory Expansion. Then, you can change the Active Memory Expansion factor, the memory assignment, or the processor assignment of the logical partition to improve its performance.

About this task

To adjust the Active Memory Expansion configuration to improve performance, complete the following steps:

Procedure

1. Run the Active Memory Expansion planning tool, which is the **amepat** command, from the AIX command line interface.

When you run the planning tool on a workload that currently uses Active Memory Expansion, the tool generates a report that provides the following information:

- Various statistics about memory compression and processor consumption.
- Several alternative configuration possibilities for Active Memory Expansion on the logical partition.
- Recommended configuration to improve the performance of Active Memory Expansion on the logical partition.

Tip: You can view more detailed statistics about memory compression and processor consumption by using the **vmstat**, **lparstat**, **svmon**, and **topas** commands.

2. Perform one or more of the following tasks to adjust the configuration:

- Dynamically change the Active Memory Expansion factor that is set for the logical partition. For instructions, see [“Changing the Active Memory Expansion factor for AIX logical partitions” on page 157](#).
- Dynamically add, move, or remove memory to or from the logical partition. For instructions, see one of the following tasks:
 - For logical partitions that use dedicated memory, see [“Managing dedicated memory dynamically” on page 156](#).

- For logical partitions that use shared memory, see [“Adding and removing logical memory dynamically to and from a shared memory partition”](#) on page 158.
- Dynamically add, move, or remove processor resources to or from the logical partition. For instructions, see [“Managing processor resources dynamically”](#) on page 160.

Performance considerations for shared memory partitions

You can learn about performance factors (such as shared memory overcommitment) that influence the performance of a logical partition that uses shared memory (also referred to as a *shared memory partition*). You can also use shared memory statistics to help you determine how to adjust the configuration of a shared memory partition to improve its performance.

Performance considerations for over committed shared memory partitions

Learn about how the degree to which the memory configuration of a logical partition that uses shared memory (also referred to as a *shared memory partition*) is over committed affects the performance of the shared memory partition. In general, the less over committed the memory configuration of a shared memory partition, the better its performance.

A shared memory configuration is considered over committed when the sum of the logical memory that is assigned to all of the shared memory partitions is greater than the amount of physical memory in the shared memory pool.

When the sum of the physical memory that is currently used by the shared memory partitions is less than or equal to the amount of memory in the shared memory pool, the memory configuration is *logically over committed*. In a logically over committed memory configuration, the shared memory pool has enough physical memory to contain the memory that is used by all shared memory partitions at one point in time.




Figure 7. A shared memory partition in a logically overcommitted memory configuration

The figure shows a shared memory partition that is assigned 2.5 GB of logical memory. Its maximum logical memory is 3 GB and its minimum logical memory is 1 GB. The figure also shows that the amount of physical memory that is currently allocated to the shared memory partition from the shared memory pool is 2.1 GB. If the workload that runs in the shared memory partition currently uses 2.1 GB of memory and requires an additional 0.2 GB of memory, and the shared memory pool is logically overcommitted, the hypervisor allocates an additional 0.2 GB of physical memory to the shared memory partition by assigning memory pages that are not currently in use by other shared memory partitions.

When the sum of the physical memory that is currently used by the shared memory partitions is greater than the amount of memory in the shared memory pool, the memory configuration is *physically over committed*. In a physically over committed memory configuration, the shared memory pool does not have enough physical memory to contain the memory that is used by all the shared memory partitions at one point in time. The hypervisor stores the difference of the physical and shared memory in the auxiliary storage.




Figure 8. A shared memory partition in a physically over committed memory configuration

The figure shows a shared memory partition that is currently allocated 0.8 GB of physical memory and assigned 2.5 GB of logical memory. If the workload that runs in the shared memory partition currently uses 0.8 GB of memory and requires an additional 1.5 GB of memory, and the shared memory pool is physically over committed, the hypervisor stores 1.5 GB of the shared memory partition's memory in its paging space device.

When the shared memory partition needs to access data on the paging space device, the hypervisor directs a paging VIOS partition to read the data from the paging space device and write the data to the shared memory pool. The more memory that the hypervisor must store on the paging space device, the more often the hypervisor and paging VIOS partition need to read and write data between the paging space device and the shared memory pool. Compared to directly accessing data that is stored in the shared memory pool, it takes more time to access data that is stored in the paging space device. Thus, in general, the less over committed the memory configuration of a shared memory partition, the better its performance.

The operating systems that run in the shared memory partitions help improve the performance of the shared memory partitions with over committed memory configurations by providing the hypervisor with information about how the operating system uses the physical memory that is allocated to it. Using this information, the hypervisor can store data that the operating system accesses the least often in the paging space device and store the data that the operating system accesses the most often in the shared memory pool. This reduces the frequency that the hypervisor needs to access the paging space device and increases the performance of the shared memory partition.

Related concepts

Factors that influence the performance of shared memory partitions

In addition to overcommitment considerations, you need to consider other factors that can affect the performance of a logical partition that uses shared memory (also referred to as a *shared memory partition*). These factors include the workload that is running in the shared memory partition, the I/O entitled memory of the shared memory partition, whether the operating system or applications that run in the shared memory partition use memory affinity, and whether the shared memory partition is configured to use redundant Virtual I/O Server (VIOS) logical partitions (also referred to as *paging VIOS partitions*).

Example: A shared memory configuration that is logically overcommitted

When the sum of the physical memory that is currently used by the shared memory partitions is less than or equal to the amount of memory in the shared memory pool, the memory configuration is *logically over committed*. In a logically over committed memory configuration, the shared memory pool has enough physical memory to contain the memory that is used by all shared memory partitions at one point in time.

Example: A shared memory configuration that is physically overcommitted

When the sum of the physical memory that is currently used by the shared memory partitions is greater than the amount of memory in the shared memory pool, the memory configuration is *physically over committed*. In a physically over committed memory configuration, the shared memory pool does not have enough physical memory to contain the memory that is used by all the shared memory partitions at one point in time. The hypervisor stores the difference of the physical and shared memory in the auxiliary storage.

Shared memory distribution

The hypervisor uses the memory weight of each logical partition that uses shared memory (hereafter referred to as *shared memory partitions*) to help determine which logical partitions receive more physical memory from the shared memory pool. To help optimize performance and memory use, the operating systems that run in shared memory partitions provide the hypervisor with information about how the operating system uses its memory to help the hypervisor determine which pages to store in the shared memory pool and which pages to store in the paging space devices.

Related reference

Performance statistics for shared memory

The Hardware Management Console (HMC) and Linux environments provide statistics about the shared memory configuration.

Factors that influence the performance of shared memory partitions

In addition to overcommitment considerations, you need to consider other factors that can affect the performance of a logical partition that uses shared memory (also referred to as a *shared memory partition*). These factors include the workload that is running in the shared memory partition, the I/O entitled memory of the shared memory partition, whether the operating system or applications that run in the shared memory partition use memory affinity, and whether the shared memory partition is configured to use redundant Virtual I/O Server (VIOS) logical partitions (also referred to as *paging VIOS partitions*).

The following table describes the types of workloads that are appropriate to run in shared memory configurations that are logically and physically over committed. It also describes the types of workloads that are not appropriate to run in a shared memory configuration.

Table 33. Workloads to run in logically over committed configurations, physically over committed configurations, and dedicated memory configurations

Workloads for logically overcommitted configurations	Workloads for physically overcommitted configurations	Workloads for dedicated memory configurations
• Workloads that peak at opposite and varying times. • Workloads with memory residency requirements that have a low average. • Workloads that do not have a sustained load. • Logical partitions that serve as failover and backup logical partitions when configured on the same server as their primary counterparts. • Test and development environments.	• Workloads that run the AIX operating system and use the file cache. • Print servers, file servers, network applications, and other workloads that are less sensitive to I/O latency. • Workloads that are inactive most of the time.	• Workloads with high quality of service criteria. • Workloads that consistently use memory resources due to sustained peak load. • High-performance computing (HPC) workloads.

In addition to the degree to which the memory configuration of a shared memory partition is over committed, the following factors can influence the performance of a shared memory partition:

- The workload that runs in a shared memory partition, the number of virtual adapters that are assigned to the shared memory partition, and the I/O entitled memory set for the shared memory partition all directly affect the performance of I/O devices. These factors can cause I/O devices to operate at their minimum memory requirements rather than their optimal memory requirements. This can cause delays in I/O operations.
- The amount of I/O entitled memory that is required for optimal performance depends on the workload and number of adapters configured.
- The operating systems that run in shared memory partitions cannot use memory affinity. Some applications rely on memory affinity to improve their performance.
- The shared memory partition might be suspended if it attempts to access data on its paging space device when the following situations occur simultaneously:
 - The paging VIOS partition becomes unavailable. For example, you shut down the paging VIOS partition or the paging VIOS partition fails.
 - The shared memory partition is not configured to use redundant paging VIOS partitions to access its paging space device.

Related concepts

Performance considerations for over committed shared memory partitions

Learn about how the degree to which the memory configuration of a logical partition that uses shared memory (also referred to as a *shared memory partition*) is over committed affects the performance of the shared memory partition. In general, the less over committed the memory configuration of a shared memory partition, the better its performance.

Related reference

Performance statistics for shared memory

The Hardware Management Console (HMC) and Linux environments provide statistics about the shared memory configuration.

Performance statistics for shared memory

The Hardware Management Console (HMC) and Linux environments provide statistics about the shared memory configuration.

Where to view statistics	Statistics to view
HMC utilization data	• Statistics about the shared memory pool, such as: – Size of the shared memory pool – Total amount of memory that is over committed – Total amount of logical memory that is assigned to the shared memory partitions – Total amount of I/O entitled memory that is assigned to the shared memory partitions – Total amount of physical memory that the shared memory partitions currently use for their I/O devices – Amount of memory from the shared memory pool that the hypervisor uses to manage the shared memory partitions – The time it takes, in microseconds, for data to be written to the shared memory pool from the paging space device • Statistics about the shared memory partitions, such as: – Amount of logical memory assigned to the shared memory partition – Amount of physical memory from the shared memory pool that is allocated to the shared memory partition – Amount of memory that is over committed – I/O entitled memory assigned to the shared memory partition – Amount of physical memory that the shared memory partition currently uses for its I/O devices – Memory weight of the shared memory partition

Where to view statistics	Statistics to view
IBM i See IBM® i to view shared memory statistics in IBM i.	• Statistics about the shared memory pool, such as: – Total number of page faults for all of the shared memory partitions – Total time, in milliseconds, that the processors waited for page faults to be resolved – Total physical memory, in bytes, that is assigned to the shared memory pool – Sum of the logical memory, in bytes, that is assigned to all of the shared memory partitions that are active – Sum of the I/O entitled memory, in bytes, that is assigned to all of the shared memory partitions that are active – Sum of the physical memory, in bytes, that the shared memory partitions that are active currently use for their I/O devices • Statistics about the shared memory partition, such as: – Memory weight of the shared memory partition – Amount of physical memory, in bytes, from the shared memory pool that is currently used by the shared memory partition – Number of times that the shared memory partition waited for a page fault – The time, in milliseconds, that the shared memory partition waited for page faults to be resolved – Maximum amount of memory, in bytes, that the shared memory partition can assign to data areas that are shared between the operating system and the server firmware – I/O entitled memory assigned to the shared memory partition – Minimum amount of physical memory, in bytes, required for all of the configured I/O devices to operate – Optimal amount of physical memory, in bytes, required for I/O devices to maximize throughput performance – Amount of physical memory, in bytes, that the shared memory partition currently uses for its I/O devices – Highest amount of physical memory, in bytes, that the shared memory partition has used for its I/O devices since the last time the shared memory partition was activated or since the last time the memory statistics were reset, whichever is most recent – Number of delayed I/O operations since the last time the shared memory partition was activated

Where to view statistics	Statistics to view
Linux View memory statistics for Linux in the <code>sysfs</code> file system as follows: - Shared memory partition data: <code>cat /proc/ppc64/lparcfg</code> - Virtual I/O bus attributes: <code>/sys/bus/vio/</code> directory. - Virtual I/O device attributes: <code>/sys/bus/vio/devices/</code> directory. This directory has a subdirectory for each device. Look in the subdirectory for each device to see the virtual I/O device statistics for each device. - Shared Memory statistics: amsstat (included in <code>powerpc-utils</code>) - Shared Memory graphical monitoring: amsvis (included in <code>powerpc-utils-python</code>)	- Statistics about the shared memory partition: - I/O entitled memory set for the shared memory partition - Memory weight of the shared memory partition - Amount of physical memory allocated to the shared memory partition - Size of the shared memory pool to which the shared memory partition belongs - Frequency that data is written to the shared memory pool from the paging space device - The time it takes, in microseconds, for data to be written to the shared memory pool from the paging space device - Statistics about the virtual I/O bus, such as the highest amount of physical memory the shared memory partition has ever used for its I/O devices. - Statistics about the virtual I/O devices, such as the frequency that the device tried to map a page to perform an I/O operation and was unable to obtain sufficient memory. In this situation, the attempt fails and delays the I/O operation. - Statistics about the tools: - The packages <code>powerpc-utils</code> and <code>powerpc-utils-python</code> are user space packages. - The amsstat script can be run from a Linux logical partition to display shared memory statistics associated with the logical partition. - The amsvis tool is a python based graphical tool that displays similar information in a graphical manner. This tool is capable of aggregating data from multiple Linux shared memory logical partitions to obtain a picture of cross logical partition performance of shared memory Linux logical partitions.

Related concepts

Factors that influence the performance of shared memory partitions

In addition to overcommitment considerations, you need to consider other factors that can affect the performance of a logical partition that uses shared memory (also referred to as a *shared memory partition*). These factors include the workload that is running in the shared memory partition, the I/O entitled memory of the shared memory partition, whether the operating system or applications that run in the shared memory partition use memory affinity, and whether the shared memory partition is configured to use redundant Virtual I/O Server (VIOS) logical partitions (also referred to as *paging VIOS partitions*).

Performance considerations for over committed shared memory partitions

Learn about how the degree to which the memory configuration of a logical partition that uses shared memory (also referred to as a *shared memory partition*) is over committed affects the performance of the shared memory partition. In general, the less over committed the memory configuration of a shared memory partition, the better its performance.

Adjusting the shared memory configuration to improve performance

You can use the Hardware Management Console (HMC) to adjust the configuration of your shared memory environment to improve its performance. For example, you can change the I/O entitled memory or the

memory weight that is assigned to a logical partition that uses shared memory (also referred to as a *shared memory partition*).

Determining the I/O entitled memory for a shared memory partition

After you create a new logical partition that uses shared memory (hereafter referred to as a *shared memory partition*) or you dynamically add or remove a virtual adapter, you can use memory statistics that are displayed by the Hardware Management Console (HMC) to dynamically increase and decrease the amount of I/O entitled memory assigned to a shared memory partition.

About this task

The I/O entitled memory set for a shared memory partition needs to be high enough to ensure the progress of I/O operations, and low enough to ensure adequate memory use among all the shared memory partitions in the shared memory pool.

The operating system manages the I/O entitled memory allocated to a shared memory partition by distributing it among the I/O device drivers. The operating system monitors how the device drivers use the I/O entitled memory, and sends usage data to the HMC. You can view the data in the HMC and dynamically adjust the I/O entitled memory that is assigned to a shared memory partition.

For more information about memory settings, see [Changing memory settings](#).

Examples

Creating a new shared memory partition

1. You activate the new shared memory partition. The HMC automatically sets the I/O entitled memory for the shared memory partition.
2. After some time, you view the memory statistics and see that the Maximum I/O Entitled Memory Used value is much less than the Assigned I/O Entitled Memory value.
3. You dynamically decrease the I/O entitled memory of the shared memory partition to the Maximum I/O Entitled Memory Used value and reset the data collector. (Dynamically decreasing the I/O entitled memory also changes the I/O entitled memory mode to the manual mode.)
4. After some time, you view the memory statistics again and determine that the new Maximum I/O Entitled Memory Used value is only slightly less than the new Assigned I/O Entitled Memory value, and no further adjustment is necessary.

Dynamically adding a virtual adapter to a shared memory partition in the auto I/O entitled memory mode

1. You dynamically add a virtual adapter to a shared memory partition. The HMC automatically increases the I/O entitled memory that is assigned to the shared memory partition.
2. After some time, you view the memory statistics and see that the Maximum I/O Entitled Memory Used value is much less than the Assigned I/O Entitled Memory value.
3. You dynamically decrease the I/O entitled memory of the shared memory partition to the Maximum I/O Entitled Memory Used value and reset the data collector. (Dynamically decreasing the I/O entitled memory also changes the I/O entitled memory mode to the manual mode.)
4. After some time, you view the memory statistics again and determine that the new Maximum I/O Entitled Memory Used value is only slightly less than the new Assigned I/O Entitled Memory value, and no further adjustment is necessary.

Dynamically adding a virtual adapter to a shared memory partition in the manual I/O entitled memory mode

1. You ensure that the shared memory partition has enough I/O entitled memory to accommodate the new adapter by dynamically increasing the I/O entitled memory of the shared memory partition.
2. You dynamically add a virtual adapter to the shared memory partition.
3. After some time, you view the memory statistics and see that the Maximum I/O Entitled Memory Used value is much less than the Assigned I/O Entitled Memory value.

4. You dynamically decrease the I/O entitled memory of the shared memory partition to the Maximum I/O Entitled Memory Used value and reset the data collector.
5. After some time, you view the memory statistics again and determine that the new Maximum I/O Entitled Memory Used value is only slightly less than the new Assigned I/O Entitled Memory value, and no further adjustment is necessary.

Dynamically removing a virtual adapter from a shared memory partition

1. You dynamically remove a virtual adapter from a shared memory partition. If the I/O entitled memory mode is in the auto mode, the HMC automatically decreases the I/O entitled memory that is assigned to the shared memory partition.
2. You reset the data collector.
3. After some time, you view the memory statistics and see that the Maximum I/O Entitled Memory Used value is much less than the Assigned I/O Entitled Memory value.
4. You dynamically decrease the I/O entitled memory of the shared memory partition to the Maximum I/O Entitled Memory Used value and reset the data collector. (If the I/O entitled memory mode is in the auto mode, dynamically decreasing the I/O entitled memory also changes the I/O entitled memory mode to the manual mode.)
5. After some time, you view the memory statistics again and determine that the new Maximum I/O Entitled Memory Used value is only slightly less than the new Assigned I/O Entitled Memory value, and no further adjustment is necessary.

The following example is another way to accomplish this example for AIX shared memory partitions:

1. Determine the amount of physical memory that the virtual adapter (that you plan to remove) currently uses by running the **lparstat** command from the AIX command line.
2. If the I/O entitled memory mode is in the auto mode, dynamically change the I/O entitled memory mode to the manual mode by running the **chhwres** command from the HMC command line.
3. Using the HMC graphical interface, dynamically remove the virtual adapter.
4. Using the HMC graphical interface, dynamically decrease the I/O entitled memory that is assigned to the shared memory partition by the amount that you identified in step “1” on [page 220](#).

Managing security for logical partitions and operating systems

When all logical partitions are managed by the Hardware Management Console, you can control who has access to the HMC and the system. You can also use the IBM eServer Security Planner to help you plan a basic security policy for each of the operating systems on your system.

When all logical partitions are managed by the Hardware Management Console (HMC), the system administrator for the HMC can control who has access to the HMC and the managed systems by creating HMC user roles. The user roles control who can access different parts of the HMC and what tasks they can perform on the managed system.

You can use the IBM eServer Security Planner to help you plan a basic security policy for each of the operating systems on your IBM Power Systems hardware. The planner provides you with a list of recommendations for setting password rules, resource-access rules, logging and auditing rules, and other security settings that are specific to the operating system.

Related information

[eServer Security Planner](#)

Troubleshooting IBM i logical partitions

If you have problems with a partitioned system, determine if the problem is specific to logical partitions or is a system problem. If your problem is specific to logical partitions, you can use the reference codes

to resolve the error. However, specific recovery actions and tasks might require the assistance of your next level of support.

Debugging network server description error messages for AIX logical partitions

This topic provides a list of network server description (NWS) error codes and explanations to help you debug NWS error messages for AIX logical partitions.

You could encounter error messages when you try to vary on an AIX logical partition. These error messages are displayed if you provide information that does not apply to a logical partition running on the server when you create your network server description (NWS). All error messages related to the NWS are displayed in QSYSOPR and indicate a description of the problem and a resolution to the problem.

Table 34. NWS error messages	
Reason codes	Code explanations
00000001	*NWSSTG was specified as the IPL source, but no storage space was found.
00000002	The partition specified in the PARTITION parameter was not found.
00000003	The partition specified in the PARTITION parameter is not a GUEST partition (that is, the TYPE parameter for the partition specified in the PARTITION parameter does not have a value of *GUEST).
00000004	There is already an NWS in the IBM i logical partition that is active and using the partition specified in the PARTITION parameter of the NWS.
00000005	The partition specified in the PARTITION parameter of the NWS is powered on (perhaps through the LPAR configuration interface or from another IBM i logical partition.)
00000006	The partition is set to start from a stream file (stmf) and that did not work. You should note that the user performing the vary on operation needs read access to the IPL STMF parameter.
00000007	The NWS is set to start from a network-storage space (NWSSTG), but the kernel could not find the NWSSTG. Some common reasons are that the storage space does not have a disk partition that is formatted as type 0x41 or is marked as startable.
00000008	The partition would not start. There are a variety of reasons why the partition will not start. You should look at the information for this partition and start reviewing the SRCs.
00000009	The partition identified as the logical partition is not configured. You should specify who has power controlling access to the partition.
00000010	A network server storage space linked to this network server is damaged. Contact your next level of support.
00000011	Contact your next level of support to find a proper solution to the problem.
00000012	The resource name you selected in the RSRNAME parameter is not valid. Use the Work with Hardware Resources (WRKHDWRSC) command with the TYPE(*CMN) parameter to help determine the resource name.
00000013	The resource you selected in the RSRNAME command exists, but is not in the partition you specified. Use the WRKHDWRSC command with the TYPE(*CMN) parameter to help determine a resource name in the partition you specified.

<i>Table 34. NWSD error messages (continued)</i>	
Reason codes	Code explanations
00000014	Unable to determine partition for resource name. Either specify a partition directly or update the resource definition at the HMC to indicate the client partition.
00000015	Unknown error occurred. Contact your next level of support.

Troubleshooting errors for Linux partitions by using IBM i virtual I/O resources

In many cases, you can troubleshoot and resolve errors specific to Linux logical partitions using IBM i virtual I/O resources without having to call service and support.

Debugging network server description error messages

This topic provides a list of network server description (NWSD) error codes and explanations to help you debug NWSD error messages for Linux logical partitions.

You could encounter error messages when you try to vary on a Linux logical partition. These error messages are displayed if you provide information when you create your network server description (NWSD) that does not apply to a logical partition running on the server. All error messages related to the NWSD are displayed in QSYSOPR indicating a description of the problem and a resolution to the problem.

<i>Table 35. NWSD error messages</i>	
Reason codes	Code explanations
00000001	*NWSSTG was specified as the IPL source, but no storage space was found.
00000002	The partition specified in the PARTITION parameter was not found. Use the CHGNWSD IBM i Control Language (CL) command to compare the partition name in the NWSD with the partition name created on the Hardware Management Console (HMC), and change the partition name as necessary.
00000003	The partition specified in the PARTITION parameter is not a GUEST partition (that is, the TYPE parameter for the partition specified in the PARTITION parameter does not have a value of *GUEST).
00000004	There is already an NWSD in the IBM i logical partition that is active and using the partition specified in the PARTITION parameter of the NWSD.
00000005	The partition specified in the PARTITION parameter of the NWSD is powered on (perhaps through the LPAR configuration interface or from another IBM i logical partition.)
00000006	The partition is set to start from a stream file (stmf) and that did not work. You should note that the user performing the vary on operation needs read access to the IPL STMF parameter.
00000007	The NWSD is set to start from a network-storage space (NWSSTG), but the kernel could not find the NWSSTG. Some common reasons are that the storage space does not have a disk partition that is formatted as type 0x41 or is marked as startable.
00000008	The partition would not start. There are a variety of reasons why the partition will not start. You should look at the information for this partition and start reviewing the SRCs.

<i>Table 35. NWSD error messages (continued)</i>	
Reason codes	Code explanations
00000009	The partition identified as the logical partition is not configured. You should specify who has power controlling access to the partition.
00000010	A network server storage space linked to this network server is damaged. Contact your next level of support.
00000011	Contact your next level of support to find a proper solution to the problem.
00000012	The resource name you selected in the RSRCTYPE parameter is not valid. Use the Work with Hardware Resources (WRKHDWRSC) command with the TYPE(*CMN) parameter to help determine the resource name.
00000013	The resource you selected in the RSRCTYPE command exists, but is not in the partition you specified. Use the WRKHDWRSC command with the TYPE(*CMN) parameter to help determine a resource name in the partition you specified.
00000014	Unable to determine partition for resource name. Either specify a partition directly or update the resource definition at the HMC to indicate the client partition.
00000015	Unknown error occurred. Contact your next level of support.

Troubleshooting Linux virtual tape errors

You can troubleshoot and recover from many common errors that are related to Linux virtual tape without having to call service and support.

If errors occur while you access Linux virtual tape, examine the file `/proc/systemi/viotape`. It describes the mapping between IBM i device names and Linux device names and records the last error for each tape device.

<i>Table 36. Common errors and recovery scenarios for troubleshooting Linux virtual tape</i>	
Error	Recovery scenario
Device unavailable	Make sure the device is varied off in the IBM i logical partition.
Not ready	Retry the operation. If the operation still fails with the same description in <code>/proc/iSeries/viotape</code> , verify that the correct medium is in the tape drive.
Load failure or cleaning cartridge found	Verify that the correct medium is in the tape drive.
Data check or Equipment check	Verify that you are using a supported block size to read or write the tape. All known tape devices that are supported by IBM can use a block size of 20 KB (supplied by the <code>-b 40</code> argument to <code>tar</code>).
Internal error	Contact your service representative.

Situations requiring the assistance of an authorized service provider

Some IBM i troubleshooting tasks on the server require the assistance of an authorized service provider. These tasks are not common and are only performed if the authorized service provider deems it necessary.

If you need to perform any of these tasks on your server, consult the [IBM Support Portal](#) website for information on server support.

Main storage dumps on IBM i logical partitions

When your system performs a main storage dump, contact service and support.

On a system with logical partitions, two types of failures can cause main storage dumps: server failure and logical partition failure.

Failures caused by server processing hardware or server firmware might cause the entire server to fail. Software failures in a logical partition cause only that logical partition to fail. A server failure may cause a platform system dump. A logical partition failure may cause a main storage dump only on that logical partition.

You can also force a main storage dump on a logical partition or managed system when you are directed to do so by an authorized service provider.

Using remote service with logical partitions

You can use the Hardware Management Console (HMC) to enable remote services with logical partitions. Remote service is a method that an authorized service provider can use to access your managed system through a modem.



Attention: Use this procedure only when directed to do so by service and support, and ensure that remote service is deactivated when your authorized service provider is finished with it. It is a security risk to leave remote service enabled when not in use. Someone could access your server without your knowledge.

1. Create a user ID.
2. Click **Service Applications** → **Remote Support** → **Customize Inbound Connectivity Settings**.

Shutting down a power domain with logical partitions

You can use the Hardware Management Console (HMC) to power off, repair, and power on the appropriate power domain when a disk unit I/O adapter (IOA) fails. By using this method, you can replace the IOA without restarting the logical partition or managed system.



Attention: Use this procedure only when directed to do so by service and support. Incorrect use of this function can cause loss of data. It can also cause failures that may be incorrectly diagnosed as expensive hardware failures.

When a disk unit IOA fails, communication with the disk units (which is controlled by the IOA) is lost, resulting in a disk unit attention SRC, and possibly partial or complete loss of system responsiveness.

Related information

[Performing dumps](#)

Troubleshooting the RMC connection between the logical partition and the HMC

To perform dynamic partitioning operations, you require a Resource Monitoring and Control (RMC) connection between the logical partition and the Hardware Management Console (HMC). If you cannot add or remove processors, memory, or I/O devices to or from a logical partition, check whether the RMC connection is active. Failure of the RMC connection is one of the most common reasons for failure of dynamic partitioning operations.

Before you begin, complete the following procedure:

1. Check the value of the RMC connection state that is cached in the data repository of the HMC by running the following command from the HMC command-line interface:

```
lssyscfg -r lpar -m cec_name -F  
name,rmc_state,rmc_ipaddr,rmc_osshutdown_capable,dlpar_mem_capable,  
dlpar_proc_capable,dlpar_io_capable
```

The value of the **rmc_state** attribute must either be active or inactive. Also, all the capabilities must be enabled.

For example:

```
#lssyscfg -r lpar -m cec_name -F
name,rmc_state,rmc_ipaddr,rmc_osshutdown_capable,dlpar_mem_capable,
dlpar_proc_capable,dlpar_io_capable
lpar01,1,9.5.23.194,1,1,1,1
...
lpar0n,1.9.5.24.###,1,1,1,1
```

If the value of the **rmc_state** attribute or all the capabilities are not set to 1, perform a system rebuild to refresh the data by running the `chsysstate -m system name -o rebuild -r sys` command. If the rebuild operation does not change the value, complete steps 2 and 3.

2. Ensure that the firewall of the HMC is lifted for the RMC port by using the HMC graphical user interface. For the procedure, see solution [1](#).
3. Ensure that the firewall of the HMC is authenticated for the HMC to receive the request from the logical partition and the logical partition is authenticated to receive the request from the HMC by either using Secure Shell (SSH) or Telnet.

When the operating system on the logical partition is Linux, ensure that Reliable Scalable Cluster Technology (RSCT) Red Hat Package Managers (RPMs) **rsct.core**, **rsct.core.utils**, and **src** are installed. For more information about how to install the RPMs, see [Service and productivity tools for SLES on POWER Linux servers for SUSE Linux Enterprise Server operating system](#) and [Service and productivity tools For Managed RHEL for Red Hat Enterprise Linux operating system](#).

The following table lists the steps to check the RMC connection and possible solutions when the connection fails.

Table 37. Steps to check for RMC failure and solutions	
Scenario	Solution
Verify whether the firewall settings block the logical partition that is managed by the HMC.	1. To verify the Firewall configuration of the LAN adapter, perform the following steps by using the HMC: a. In the navigation pane, open HMC Management. b. In the work pane, click Change Network Settings. c. Click the LAN Adapters tab. d. Select any LAN adapter other than the eth0 adapter that connects the HMC with the service processor, and click Details. e. On the LAN Adapter tab, under Local area network information, verify whether Open is selected and Partition communication status is displayed as enabled. f. Click the Firewall Settings tab. g. Ensure that the RMC application is one of the applications that are displayed in Allowed Hosts. If it is not displayed in Allowed Hosts, select the RMC application under Available Applications and click Allow Incoming. h. Click OK.

Table 37. Steps to check for RMC failure and solutions (continued)

Scenario	Solution
Verify whether the /tmp folder in the HMC is 100% full by running the df command, with superuser privilege.	You must remove unused files in the /tmp folder to free up space.

Related information

[Checking the status of the management domain and the peer domain](#)

[Verifying RMC connections for the mobile partition](#)

[RMC network port usage, data flows, and security](#)

Performance considerations for logical partitions

You can manage and enhance the performance of logical partitions so that your system uses its resources in the most efficient manner.

You can manage and enhance the performance of a AIX logical partition by configuring the AIX operating system.

Managing IBM i performance ensures that your managed system is efficiently using resources and that your managed system provides the best possible services to you and to your business. Moreover, effective performance management can help you quickly respond to changes in your managed system and can save on expenses by postponing costly upgrades and service fees.

Related information

[AIX Performance management](#)

[Performance Tools Guide and Reference](#)

[Performance Toolbox Version 2 and 3 Guide and Reference](#)

[Power Systems Capacity on Demand](#)

Performance considerations for shared memory partitions

You can learn about performance factors (such as shared memory overcommitment) that influence the performance of a logical partition that uses shared memory (also referred to as a *shared memory partition*). You can also use shared memory statistics to help you determine how to adjust the configuration of a shared memory partition to improve its performance.

Performance considerations for over committed shared memory partitions

Learn about how the degree to which the memory configuration of a logical partition that uses shared memory (also referred to as a *shared memory partition*) is over committed affects the performance of the shared memory partition. In general, the less over committed the memory configuration of a shared memory partition, the better its performance.

A shared memory configuration is considered over committed when the sum of the logical memory that is assigned to all of the shared memory partitions is greater than the amount of physical memory in the shared memory pool.

When the sum of the physical memory that is currently used by the shared memory partitions is less than or equal to the amount of memory in the shared memory pool, the memory configuration is *logically over committed*. In a logically over committed memory configuration, the shared memory pool has enough physical memory to contain the memory that is used by all shared memory partitions at one point in time.




Figure 9. A shared memory partition in a logically overcommitted memory configuration

The figure shows a shared memory partition that is assigned 2.5 GB of logical memory. Its maximum logical memory is 3 GB and its minimum logical memory is 1 GB. The figure also shows that the amount of physical memory that is currently allocated to the shared memory partition from the shared memory pool is 2.1 GB. If the workload that runs in the shared memory partition currently uses 2.1 GB of memory and requires an additional 0.2 GB of memory, and the shared memory pool is logically overcommitted, the hypervisor allocates an additional 0.2 GB of physical memory to the shared memory partition by assigning memory pages that are not currently in use by other shared memory partitions.

When the sum of the physical memory that is currently used by the shared memory partitions is greater than the amount of memory in the shared memory pool, the memory configuration is *physically over committed*. In a physically over committed memory configuration, the shared memory pool does not have enough physical memory to contain the memory that is used by all the shared memory partitions at one point in time. The hypervisor stores the difference of the physical and shared memory in the auxiliary storage.




Figure 10. A shared memory partition in a physically over committed memory configuration

The figure shows a shared memory partition that is currently allocated 0.8 GB of physical memory and assigned 2.5 GB of logical memory. If the workload that runs in the shared memory partition currently uses 0.8 GB of memory and requires an additional 1.5 GB of memory, and the shared memory pool is physically over committed, the hypervisor stores 1.5 GB of the shared memory partition's memory in its paging space device.

When the shared memory partition needs to access data on the paging space device, the hypervisor directs a paging VIOS partition to read the data from the paging space device and write the data to the shared memory pool. The more memory that the hypervisor must store on the paging space device, the more often the hypervisor and paging VIOS partition need to read and write data between the paging space device and the shared memory pool. Compared to directly accessing data that is stored in the shared memory pool, it takes more time to access data that is stored in the paging space device. Thus, in general, the less over committed the memory configuration of a shared memory partition, the better its performance.

The operating systems that run in the shared memory partitions help improve the performance of the shared memory partitions with over committed memory configurations by providing the hypervisor with information about how the operating system uses the physical memory that is allocated to it. Using this information, the hypervisor can store data that the operating system accesses the least often in the paging space device and store the data that the operating system accesses the most often in the shared memory pool. This reduces the frequency that the hypervisor needs to access the paging space device and increases the performance of the shared memory partition.

Related concepts

Factors that influence the performance of shared memory partitions

In addition to overcommitment considerations, you need to consider other factors that can affect the performance of a logical partition that uses shared memory (also referred to as a *shared memory partition*). These factors include the workload that is running in the shared memory partition, the I/O entitled memory of the shared memory partition, whether the operating system or applications that run in the shared memory partition use memory affinity, and whether the shared memory partition is configured to use redundant Virtual I/O Server (VIOS) logical partitions (also referred to as *paging VIOS partitions*).

Example: A shared memory configuration that is logically overcommitted

When the sum of the physical memory that is currently used by the shared memory partitions is less than or equal to the amount of memory in the shared memory pool, the memory configuration is *logically over committed*. In a logically over committed memory configuration, the shared memory pool has enough physical memory to contain the memory that is used by all shared memory partitions at one point in time.

Example: A shared memory configuration that is physically overcommitted

When the sum of the physical memory that is currently used by the shared memory partitions is greater than the amount of memory in the shared memory pool, the memory configuration is *physically over committed*. In a physically over committed memory configuration, the shared memory pool does not have enough physical memory to contain the memory that is used by all the shared memory partitions at one point in time. The hypervisor stores the difference of the physical and shared memory in the auxiliary storage.

Shared memory distribution

The hypervisor uses the memory weight of each logical partition that uses shared memory (hereafter referred to as *shared memory partitions*) to help determine which logical partitions receive more physical memory from the shared memory pool. To help optimize performance and memory use, the operating systems that run in shared memory partitions provide the hypervisor with information about how the operating system uses its memory to help the hypervisor determine which pages to store in the shared memory pool and which pages to store in the paging space devices.

Related reference

Performance statistics for shared memory

The Hardware Management Console (HMC) and Linux environments provide statistics about the shared memory configuration.

Factors that influence the performance of shared memory partitions

In addition to overcommitment considerations, you need to consider other factors that can affect the performance of a logical partition that uses shared memory (also referred to as a *shared memory partition*). These factors include the workload that is running in the shared memory partition, the I/O entitled memory of the shared memory partition, whether the operating system or applications that run in the shared memory partition use memory affinity, and whether the shared memory partition is configured to use redundant Virtual I/O Server (VIOS) logical partitions (also referred to as *paging VIOS partitions*).

The following table describes the types of workloads that are appropriate to run in shared memory configurations that are logically and physically over committed. It also describes the types of workloads that are not appropriate to run in a shared memory configuration.

Table 38. Workloads to run in logically over committed configurations, physically over committed configurations, and dedicated memory configurations

Workloads for logically overcommitted configurations	Workloads for physically overcommitted configurations	Workloads for dedicated memory configurations
• Workloads that peak at opposite and varying times. • Workloads with memory residency requirements that have a low average. • Workloads that do not have a sustained load. • Logical partitions that serve as failover and backup logical partitions when configured on the same server as their primary counterparts. • Test and development environments.	• Workloads that run the AIX operating system and use the file cache. • Print servers, file servers, network applications, and other workloads that are less sensitive to I/O latency. • Workloads that are inactive most of the time.	• Workloads with high quality of service criteria. • Workloads that consistently use memory resources due to sustained peak load. • High-performance computing (HPC) workloads.

In addition to the degree to which the memory configuration of a shared memory partition is over committed, the following factors can influence the performance of a shared memory partition:

- The workload that runs in a shared memory partition, the number of virtual adapters that are assigned to the shared memory partition, and the I/O entitled memory set for the shared memory partition all directly affect the performance of I/O devices. These factors can cause I/O devices to operate at their minimum memory requirements rather than their optimal memory requirements. This can cause delays in I/O operations.
- The amount of I/O entitled memory that is required for optimal performance depends on the workload and number of adapters configured.
- The operating systems that run in shared memory partitions cannot use memory affinity. Some applications rely on memory affinity to improve their performance.
- The shared memory partition might be suspended if it attempts to access data on its paging space device when the following situations occur simultaneously:
 - The paging VIOS partition becomes unavailable. For example, you shut down the paging VIOS partition or the paging VIOS partition fails.
 - The shared memory partition is not configured to use redundant paging VIOS partitions to access its paging space device.

Related concepts

Performance considerations for over committed shared memory partitions

Learn about how the degree to which the memory configuration of a logical partition that uses shared memory (also referred to as a *shared memory partition*) is over committed affects the performance of the shared memory partition. In general, the less over committed the memory configuration of a shared memory partition, the better its performance.

Related reference

Performance statistics for shared memory

The Hardware Management Console (HMC) and Linux environments provide statistics about the shared memory configuration.

Performance statistics for shared memory

The Hardware Management Console (HMC) and Linux environments provide statistics about the shared memory configuration.

Where to view statistics	Statistics to view
<u>HMC utilization data</u>	• Statistics about the shared memory pool, such as: – Size of the shared memory pool – Total amount of memory that is over committed – Total amount of logical memory that is assigned to the shared memory partitions – Total amount of I/O entitled memory that is assigned to the shared memory partitions – Total amount of physical memory that the shared memory partitions currently use for their I/O devices – Amount of memory from the shared memory pool that the hypervisor uses to manage the shared memory partitions – The time it takes, in microseconds, for data to be written to the shared memory pool from the paging space device • Statistics about the shared memory partitions, such as: – Amount of logical memory assigned to the shared memory partition – Amount of physical memory from the shared memory pool that is allocated to the shared memory partition – Amount of memory that is over committed – I/O entitled memory assigned to the shared memory partition – Amount of physical memory that the shared memory partition currently uses for its I/O devices – Memory weight of the shared memory partition

Where to view statistics	Statistics to view
IBM i See IBM® i to view shared memory statistics in IBM i.	• Statistics about the shared memory pool, such as: – Total number of page faults for all of the shared memory partitions – Total time, in milliseconds, that the processors waited for page faults to be resolved – Total physical memory, in bytes, that is assigned to the shared memory pool – Sum of the logical memory, in bytes, that is assigned to all of the shared memory partitions that are active – Sum of the I/O entitled memory, in bytes, that is assigned to all of the shared memory partitions that are active – Sum of the physical memory, in bytes, that the shared memory partitions that are active currently use for their I/O devices • Statistics about the shared memory partition, such as: – Memory weight of the shared memory partition – Amount of physical memory, in bytes, from the shared memory pool that is currently used by the shared memory partition – Number of times that the shared memory partition waited for a page fault – The time, in milliseconds, that the shared memory partition waited for page faults to be resolved – Maximum amount of memory, in bytes, that the shared memory partition can assign to data areas that are shared between the operating system and the server firmware – I/O entitled memory assigned to the shared memory partition – Minimum amount of physical memory, in bytes, required for all of the configured I/O devices to operate – Optimal amount of physical memory, in bytes, required for I/O devices to maximize throughput performance – Amount of physical memory, in bytes, that the shared memory partition currently uses for its I/O devices – Highest amount of physical memory, in bytes, that the shared memory partition has used for its I/O devices since the last time the shared memory partition was activated or since the last time the memory statistics were reset, whichever is most recent – Number of delayed I/O operations since the last time the shared memory partition was activated

Where to view statistics	Statistics to view
Linux View memory statistics for Linux in the <code>sysfs</code> file system as follows: • Shared memory partition data: <code>cat /proc/ppc64/lparcfg</code> • Virtual I/O bus attributes: <code>/sys/bus/vio/</code> directory. • Virtual I/O device attributes: <code>/sys/bus/vio/devices/</code> directory. This directory has a subdirectory for each device. Look in the subdirectory for each device to see the virtual I/O device statistics for each device. • Shared Memory statistics: amsstat (included in <code>powerpc-utils</code>) • Shared Memory graphical monitoring: amsvis (included in <code>powerpc-utils-python</code>)	• Statistics about the shared memory partition: – I/O entitled memory set for the shared memory partition – Memory weight of the shared memory partition – Amount of physical memory allocated to the shared memory partition – Size of the shared memory pool to which the shared memory partition belongs – Frequency that data is written to the shared memory pool from the paging space device – The time it takes, in microseconds, for data to be written to the shared memory pool from the paging space device • Statistics about the virtual I/O bus, such as the highest amount of physical memory the shared memory partition has ever used for its I/O devices. • Statistics about the virtual I/O devices, such as the frequency that the device tried to map a page to perform an I/O operation and was unable to obtain sufficient memory. In this situation, the attempt fails and delays the I/O operation. • Statistics about the tools: – The packages <code>powerpc-utils</code> and <code>powerpc-utils-python</code> are user space packages. – The amsstat script can be run from a Linux logical partition to display shared memory statistics associated with the logical partition. – The amsvis tool is a python based graphical tool that displays similar information in a graphical manner. This tool is capable of aggregating data from multiple Linux shared memory logical partitions to obtain a picture of cross logical partition performance of shared memory Linux logical partitions.

Related concepts

Factors that influence the performance of shared memory partitions

In addition to overcommitment considerations, you need to consider other factors that can affect the performance of a logical partition that uses shared memory (also referred to as a *shared memory partition*). These factors include the workload that is running in the shared memory partition, the I/O entitled memory of the shared memory partition, whether the operating system or applications that run in the shared memory partition use memory affinity, and whether the shared memory partition is configured to use redundant Virtual I/O Server (VIOS) logical partitions (also referred to as *paging VIOS partitions*).

Performance considerations for over committed shared memory partitions

Learn about how the degree to which the memory configuration of a logical partition that uses shared memory (also referred to as a *shared memory partition*) is over committed affects the performance of the shared memory partition. In general, the less over committed the memory configuration of a shared memory partition, the better its performance.

Adjusting the shared memory configuration to improve performance

You can use the Hardware Management Console (HMC) to adjust the configuration of your shared memory environment to improve its performance. For example, you can change the I/O entitled memory or the

memory weight that is assigned to a logical partition that uses shared memory (also referred to as a *shared memory partition*).

Troubleshooting the RMC connection between the logical partition and the HMC

To perform dynamic partitioning operations, you require a Resource Monitoring and Control (RMC) connection between the logical partition and the Hardware Management Console (HMC). If you cannot add or remove processors, memory, or I/O devices to or from a logical partition, check whether the RMC connection is active. Failure of the RMC connection is one of the most common reasons for failure of dynamic partitioning operations.

Before you begin, complete the following procedure:

1. Check the value of the RMC connection state that is cached in the data repository of the HMC by running the following command from the HMC command-line interface:

```
lssyscfg -r lpar -m cec_name -F
name,rmc_state,rmc_ipaddr,rmc_osshutdown_capable,dlpar_mem_capable,
dlpar_proc_capable,dlpar_io_capable
```

The value of the **rmc_state** attribute must either be active or inactive. Also, all the capabilities must be enabled.

For example:

```
#lssyscfg -r lpar -m cec_name -F
name,rmc_state,rmc_ipaddr,rmc_osshutdown_capable,dlpar_mem_capable,
dlpar_proc_capable,dlpar_io_capable
lpar01,1,9.5.23.194,1,1,1,1
...
lpar0n,1.9.5.24.###,1,1,1,1
```

If the value of the **rmc_state** attribute or all the capabilities are not set to 1, perform a system rebuild to refresh the data by running the `chsysstate -m system name -o rebuild -r sys` command. If the rebuild operation does not change the value, complete steps [2](#) and [3](#).

2. Ensure that the firewall of the HMC is lifted for the RMC port by using the HMC graphical user interface. For the procedure, see [solution 1](#).
3. Ensure that the firewall of the HMC is authenticated for the HMC to receive the request from the logical partition and the logical partition is authenticated to receive the request from the HMC by either using Secure Shell (SSH) or Telnet.

When the operating system on the logical partition is Linux, ensure that Reliable Scalable Cluster Technology (RSCT) Red Hat Package Managers (RPMs) **rsct.core**, **rsct.core.utils**, and **src** are installed. For more information about how to install the RPMs, see [Service and productivity tools for SLES on POWER Linux servers](#) for SUSE Linux Enterprise Server operating system and [Service and productivity tools For Managed RHEL](#) for Red Hat Enterprise Linux operating system.

The following table lists the steps to check the RMC connection and possible solutions when the connection fails.

Table 39. Steps to check for RMC failure and solutions

Scenario	Solution
Verify whether the firewall settings block the logical partition that is managed by the HMC.	- To verify the Firewall configuration of the LAN adapter, perform the following steps by using the HMC: - In the navigation pane, open HMC Management. - In the work pane, click Change Network Settings. - Click the LAN Adapters tab. - Select any LAN adapter other than the eth0 adapter that connects the HMC with the service processor, and click Details. - On the LAN Adapter tab, under Local area network information, verify whether Open is selected and Partition communication status is displayed as enabled. - Click the Firewall Settings tab. - Ensure that the RMC application is one of the applications that are displayed in Allowed Hosts. If it is not displayed in Allowed Hosts, select the RMC application under Available Applications and click Allow Incoming. - Click OK.
Verify whether the /tmp folder in the HMC is 100% full by running the df command, with superuser privilege.	You must remove unused files in the /tmp folder to free up space.

Related information

[Checking the status of the management domain and the peer domain](#)

[Verifying RMC connections for the mobile partition](#)

[RMC network port usage, data flows, and security](#)

Notices

This information was developed for products and services offered in the US.

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, MD-NC119
Armonk, NY 10504-1785
US

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.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

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 jurisdictions 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide 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 Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

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 or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

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 actual people or business enterprises 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.

Accessibility features for IBM Power Systems servers

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Overview

The IBM Power Systems servers include the following major accessibility features:

- Keyboard-only operation
- Operations that use a screen reader

The IBM Power Systems servers use the latest W3C Standard, [WAI-ARIA 1.0 \(www.w3.org/TR/wai-aria/\)](http://www.w3.org/TR/wai-aria/), to ensure compliance with US Section 508 (www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards) and [Web Content](#)

Accessibility Guidelines (WCAG) 2.0 (www.w3.org/TR/WCAG20/). To take advantage of accessibility features, use the latest release of your screen reader and the latest web browser that is supported by the IBM Power Systems servers.

The IBM Power Systems servers online product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described in the Accessibility section of the IBM Knowledge Center help (www.ibm.com/support/knowledgecenter/doc/kc_help.html#accessibility).

Keyboard navigation

This product uses standard navigation keys.

Interface information

The IBM Power Systems servers user interfaces do not have content that flashes 2 - 55 times per second.

The IBM Power Systems servers web user interface relies on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

The IBM Power Systems servers web user interface includes WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Vendor software

The IBM Power Systems servers include certain vendor software that is not covered under the IBM license agreement. IBM makes no representation about the accessibility features of these products. Contact the vendor for accessibility information about its products.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service
800-IBM-3383 (800-426-3383)
(within North America)

For more information about the commitment that IBM has to accessibility, see [IBM Accessibility \(www.ibm.com/able\)](http://www.ibm.com/able).

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as the customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details/us/en/> in the section entitled "Cookies, Web Beacons and Other Technologies".

Programming interface information

Logical partitioning publication documents intended Programming Interfaces that allow the customer to write programs to obtain the services of IBM AIX Version 7.2, IBM AIX Version 7.1, IBM AIX Version 6.1, IBM i 7.4, and IBM Virtual I/O Server Version 3.1.2.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.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](#).

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

Red Hat, JBoss, OpenShift, Fedora, Hibernate, Ansible, CloudForms, RHCA, RHCE, RHCSA, Ceph, and Gluster are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group 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.

