

IBM

@server

iSeries

Concepts





@server

iSeries

Concepts

Contents

Logical partition concepts	1
How logical partitioning works	1
How logical partitioning can work for you	3
Hardware for logical partitions	4
Logical partition concept: bus	5
Logical partition concept: Bus-level and IOP-level I/O partitions	6
Dynamically switching IOPs between partitions	7
Logical partition concept: IOP	8
Logical partition concept: SPD and PCI	10
Logical partition concept: Processor.	10
Logical partition concept: memory	12
Logical partition concept: disk units	13
Logical partition concept: interactive performance.	15
Software licensing and licensed programs for logical partitions	15
Logical partition release support	16
OS/400 logical partition function by release	17
Communication options for logical partitions.	18

Logical partition concepts

The iSeries server offers you the ability to partition one server into several independent servers. Before you start creating partitions, it is essential that you understand the concepts behind this type of system configuration. The purpose of this topic is to familiarize you with the hardware and software required for logical partitions and to prepare you to plan and create logical partitions on your iSeries.

“How logical partitioning works”

Understand the meaning of a logical partitioned system and how primary and secondary partitions operate as independent servers.

“How logical partitioning can work for you” on page 3

Understand the advantages of partitioning your server and practical scenarios your company can use with this advanced technology.

“Hardware for logical partitions” on page 4

Understand the basic hardware concepts and requirements to partition your server.


“Software licensing and licensed programs for logical partitions” on page 15

Learn about the software licensing and pricing strategy for IBM products on an iSeries server with logical partitions.

“Communication options for logical partitions” on page 18

Learn how logical partitions are able to share data between partitions or between servers.

How logical partitioning works



Logical partitioning is the ability to make an iSeries server run as if it were two or more independent servers. Each logical partition operates as an independent logical server. However, each partition shares a few physical system attributes such as the system serial number, system model, and processor feature code. All other system attributes may vary among partitions.

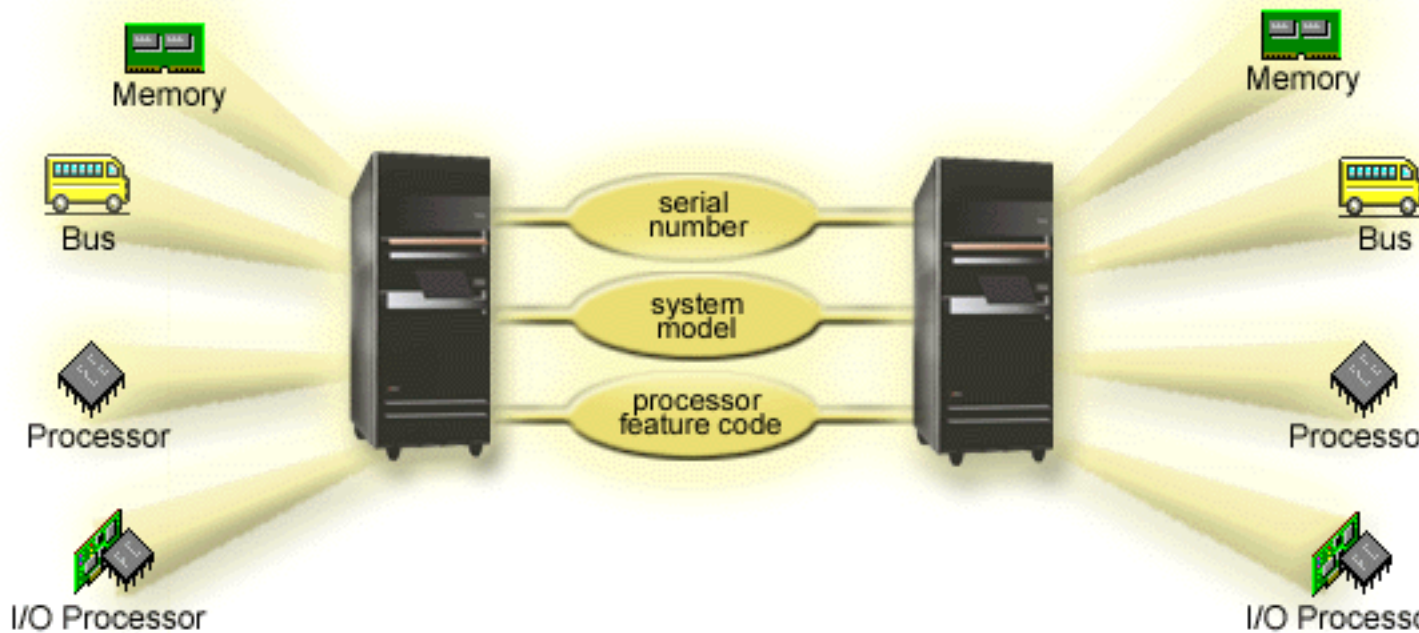
Logical partitions fall into two categories, primary partitions or secondary partitions. Each logically partitioned system has one primary partition and one or more secondary partitions. The primary partition is the only partition prior to performing any configuration changes on a system. Before secondary partitions are created, all system resources are assigned to the primary partition. Secondary partitions are independent of each other. While each secondary partition maintains a dependency on the primary, it otherwise operates as a stand-alone server.

All management functions for partitions are integrated into the Licensed Internal Code of the primary partition.

When you perform a restart of the primary partition on a server with multiple partitions, the primary partition starts first. The primary partition owns certain system resources (the system operation panel, service processor, and system key lock). The primary partition can start (IPL) secondary partitions once the system validates these resources. If a processor, memory card, or system bus failure occurs, you can find system error log entries in the primary partition Product Activity Log.

The primary partition, essentially the partition manager for your server, must remain active for the secondary partitions to be active. It is important to plan carefully on how you operate the primary partition or the types of workload you run in the primary partition. For example, OS/400 commands such as Power Down System (PWRDWNSYS), operator panel functions such as 3, 8, or 22, applying fixes (PTFs) that require a restart will affect all of the secondary partitions. You may want to restrict the primary partition to simple partition management tasks only. Since all logical partition resource movement is achieved through use of the primary partition, the isolation of the primary partition provides a secure environment whereby users in secondary partitions are not able to move resources such as processor or memory without going through a logical partition administrator of the primary partition. Where primary partition isolation is not possible, you may want to think about implementing applications that require little or no maintenance, by not using it as a test partition.

Each logical partition represents a division of resources in your iSeries server. Each partition is logical because the division of resources is virtual, not along physical boundaries. The primary resources in your server are its processors, memory, buses, and IOPs. The following diagram shows the division of system resources on a server that has two partitions:



How logical partitioning can work for you

Logical partitions on an iSeries server could prove beneficial in the following scenarios:

Creating a Mixed Production and Test Environment

You can create a combination production and test environment on the same server. You can use a logical partition as a test partition or a production partition. A production partition runs your main business applications. A failure in a production partition could significantly hinder business operations and cost the customer time and money. A test partition tests software. This might include OS/400 release testing. A failure in a test partition, while not necessarily planned, will not disrupt normal business operations.

Creating a Multiple Production Partition Environment

You create multiple production partitions in your secondary partitions. In this situation, you dedicate the primary partition to partition management.

Consolidation

A logically partitioned server can reduce the number of servers that are needed within an enterprise. You can consolidate several servers into a single logically partitioned system. This eliminates the need for, and expense of, additional equipment. You can shift resources from one logical partition to another as needs change.

Hot Backup

When a secondary partition replicates another logical partition within the same system, to switch to the backup during partition failure would cause minimal inconvenience. This configuration also minimizes the

effect of long save windows. You can take the backup partition off line and save, while the other logical partition continues to perform production work. You will need special software to use this hot backup strategy.

Integrated Cluster

Using OptiConnect and 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 secondary partition.

Maintaining Independent Systems

Dedicating a portion of the resources (disk storage unit, processors, memory, and I/O devices) to a partition achieves logical isolation of software. Logical partitions also have some hardware fault tolerance if configured properly. Interactive and batch workloads which may not run well together on a single machine can be isolated and run efficiently in separate partitions.

Running Linux

You could consolidate several Linux servers to an iSeries server. Linux enhances iSeries flexibility by enabling another application environment. Linux applications can access DB2 UDB as well as OS/400 programs and services using virtual Ethernet.

For a better understanding of how your company might use logical and guest partition, read Logical and guest partition scenarios.

Hardware for logical partitions


The information in this topic describes the required hardware needed on your server to successfully create a logical partition.

Attention:

The ability to dynamically move resources between partitions is available on both AS/400 and iSeries models that support logical partitions.

However, uni-processor partitioning capability and shared processor pool capability is only available on V5R1 and V5R2 partitions running on iSeries 820s, 830s, 840s, and some 270s. Other hardware for the logical partition is either optional, or you can switch it between two or more logical partitions.

You can read Plan for logical partitions to help you decide what hardware you need. To learn more about the hardware you can use with logical partitions, click on the piece of hardware you want to learn more about in the following picture:



To display system hardware resources, follow these steps:

1. In iSeries Navigator, expand **My Connections** or your active environment.
2. Select the primary partition of the system.
3. Expand **Configuration and Service** and select **Logical Partitions**.
4. Right-click the **Logical Partition** and select **Configure Partitions**. You are now working in the Configure Logical Partitions window.
5. Select **Physical System** to view the hardware resources for the entire system.

Logical partition concept: bus

A bus is a conductor that is used for transmitting signals or power.

A system I/O bus conducts instructions from “Logical partition concept: memory” on page 12 to the devices that are attached to the “Logical partition concept: IOP” on page 8. A system I/O bus also conducts instructions from the IOPs back to memory.

The main system unit contains one system I/O bus that the primary partition always uses bus 1. Secondary partitions can share this bus. Most “Logical partition concept: expansion unit” on page 14 contain at least one bus as well.

Every logical partition needs a bus which the logical partition can either own dedicated or shared with other logical partitions. Each logical partition can use (without owning) a system I/O bus. You can dynamically change which logical partition owns the bus or the type of bus ownership (shared or dedicated) provided the owning partition owns the bus shared.

It is possible to divide up your resources by system I/O bus when you create your logical partitions. This is called bus-level I/O partitioning. In this situation, you assign all resources (IOPs, IOAs, and devices) that are attached to the bus to only one logical partition.

It is also possible to share a bus and divide up the resources on the bus according to IOPs. This is called IOP-level I/O partitioning. In this situation, you assign all resources (IOAs and devices) that are attached to a single IOP to only one logical partition at a time. You may assign other IOPs attached to this same bus to any other (or to the same) logical partition.

When you add a bus to a logical partition, you must select whether to share the bus with other logical partitions. You have the following options for bus ownership type:

- Own dedicated: you assign all IOPs, resources, and empty card slots to the partition (bus-level I/O partitioning).
- Own bus shared: Some IOPs along with their empty card slots can be assigned go with the bus owner (IOP-level I/O partitioning).
- Use bus shared: Another logical partition lists the bus as own bus shared, but this logical partition also uses the bus. In this situation, the owning logical partition may experience slightly improved performance. Consider having the partition using high data transfer rate resources to own the bus.

You must assign the *own dedicated* type to any bus that contains the OptiConnect hardware.

From the Configure Logical partition window, you can view all of the buses on the system. You may perform concurrent maintenance from the owning logical partition (dedicated buses only) or from the primary partition. However, you must perform concurrent maintenance for shared buses from the primary partition.

“Hardware for logical partitions” on page 4

Logical partition concept: Bus-level and IOP-level I/O partitions

Depending on your needs, there can be advantages to setting up one type of I/O partition instead of another.

With bus level I/O partitions, the system partitions I/O resources by “Logical partition concept: bus” on page 5. On a server that is completely partitioned at the bus level, every secondary partition has its own removable media and workstation.

Bus-level logical partitions allow for:

- Better problem isolation and therefore higher availability.
- Better performance.
- Simplified hardware management.

When you partition a server at the “Logical partition concept: IOP” on page 8 level, one or more buses are shared and divided up between the I/O resources by the IOP. This type of logical partitions allows for:

- Greater flexibility when partitioning I/O subsystem.
- Potential cost reduction by eliminating some expansion units which the server requires to support additional buses.
- Optimization of hardware resources to avoid server limits.
- The ability to dynamically move the control of an IOP from one partition to another without the need to restart the system.
- Simplified configuration planning since you can dynamically move hardware from one partition to another.

You might also consider having a system configuration with both bus-level partitioning and IOP-level partitioning. For example, you may put all of the IOPs that you wish to switch in a shared bus and configure all other logical partitions to have bus-level partitioning. The shared bus can then belong to a test partition. This allows you the ability to switch IOPs such as tape drives or LAN adapters to partitions that need these resources.

Dynamically switching IOPs between partitions

A powerful advantage of logical partitions is ability to dynamically switch an “Logical partition concept: IOP” on page 8 from one partition to another. In other words, you can take the control of an IOP away from one partition and give it to another without restarting the server.

On a shared bus, IOPs and their resources (all IOAs and attached devices) can be dynamically moved between partitions. To initiate this action, right click on the desired IOP and select **Move**.

When you dynamically switch an IOP between partitions, you make it appear as though partitions can share devices. For example, if your server has a low use device, more than one partition can use that device by switching IOPs. Each partition can use the device because switching an IOP switches all devices attached to it. However, partitions can only use that device one at a time. You can only switch the IOP to one partition at a time. Before you perform this task, you need to free that device in OS/400 on the source partition. You can do this by using the iSeries Navigator.

IOPs that are candidates for switching include:

- IOPs that control high cost devices.
- IOPs that control low use devices and low demand devices.
- IOPs that solely control the target device or devices.

Before you use IOP switching, you should remember the IBM recommended approach to logical partitioning is to dedicate all hardware within a partition. You should partition your server at the “Logical partition concept: Bus-level and IOP-level I/O partitions” on page 6. However, when you use bus level partitioning you cannot switch IOPs. Also, bus level partitioning is not always cost-effective. Therefore, you may find it preferable to share some devices among partitions. And while partitions cannot concurrently share devices, IOP switching may be an effective solution to the cost problem.

Before you implement switching, you should consider other possible alternatives. To share devices among partitions, you can apply the same techniques that are used to share devices among separate physical systems:

- Use multiple IOPs, one in each partition, for devices that support multiple connections (some high end tape drives).
- Use multiple IOPs, one in each partition, and a switch box for devices that only support single connections (printers, or some high end tape drives).
- Use multiple IOPs and multiple devices in each partition for a self-contained solution (internal removable media devices).
- If it is not feasible for you to use any of these alternatives, then you should implement IOP switching.

You can implement dynamic IOP switching by making sure that the primary partition owns each bus with a “Logical partition concept: bus” on page 5 of *own bus shared*. All secondary partitions can then simply use any bus they need. With this configuration you can assign every IOP in the server to the partition that wants to use it. All partitions can dynamically add and remove any IOPs that control high cost, and low use devices that are attached to the system.

In addition to enabling IOP switching, this configuration has several other advantages.

- Conceptually, the implementation is easy to configure and understand.
- The primary partition will own any new hardware that is added to the system.

- As partition needs change over time, you can make adjustments with optimal hardware use and greater flexibility.

Logical partition concept: IOP

The IOP attaches to the “Logical partition concept: bus” on page 5 and one or more input/output adapters (IOAs). The IOP processes instructions from the server and works with the IOAs to control the I/O devices.

There are many different kinds of IOPs.

- Some IOPs can only support one type of I/O device. In this case the IOA is embedded in the IOP so you can not remove the IOA or change it.
- Some IOPs can support multiple device types, but only one at a time. The type of IOA that is attached determines what device you can use. IOAs on these IOPs can be changed with another IOA to support a different I/O device. The IOA and the IOP work together to control the device.
- Some IOPs can support multiple types of I/O devices at the same time. These are known as Multifunction IOPs (MFIOP) or Combined function IOPs (CFIOP). Multifunction IOPs connect to a variety of different IOAs. For instance, an MFIOP could support disk units, workstations, a communications line, and removeable media devices.

A Combined function IOP can connect to a variety of different IOAs. For instance, a CFIOP could support disk units, a console, and communications hardware. The CFIOP contains some of the same capabilities as an MFIOP, as well as Ethernet and token-ring controllers. IOAs for the supported types of I/O devices attach to the IOP.

There are several important I/O devices in the server. These include the load source disk unit, the alternate IPL device, the system console, and the electronic customer support hardware. The server needs to know where to locate these special devices on secondary partitions. When you create a logical partition, you need to identify the IOPs that control these important devices:

- The IOP that controls the “Logical partition concept: disk units” on page 13 that will be the “Logical partition concept: load source” on page 14.
- The IOP that controls the “Logical partition concept: console” on page 14.
- The IOP that controls the “Logical partition concept: alternate restart (IPL) device and removable media device” on page 13.
- The IOP that controls the electronic customer support line.

Note:

A system with logical partitions needs to have the correct IOP feature codes for the load source disk unit and alternate IPL devices. Without the correct hardware, the secondary partitions will not function correctly.

A logical partition controls all devices connected to an IOP. You cannot switch one I/O device to another logical partition without moving the ownership of IOP.

On a shared bus, you can move IOPs from a logical partition without restarting the server.

To move an IOP from a logical partition, all resources (IOAs and devices) that are attached to the IOP cannot be in use.

To move an I/O processor from one logical partition to another partition, follow these steps:

1. Ensure that any devices attached to the I/O processor you want to move are not busy. The devices should be varied off and listed as unavailable hardware.
2. In iSeries Navigator, expand **My Connections** or your active environment.
3. Select the primary partition of the system.
4. Expand **Configuration and Service** and select **Logical Partitions**.

5. Right-click the **Logical Partition** and select **Configure Partitions**. You are now working in the Configure Logical Partitions window.
6. Select the partition that has the I/O processor you would like to move.
7. Right-click the I/O processor you want move, and select **Move**.

The system reports all errors that involve IOPs to the product activity log (PAL) in the logical partition that owns the IOP. However, the system may report errors in the PAL of the primary partition against load source IOPs. This may occur when a secondary partition restarts.

“Hardware for logical partitions” on page 4

Logical partition concept: IOP and device switching

If you elect to partition at the “Logical partition concept: Bus-level and IOP-level I/O partitions” on page 6, you can switch certain IOPs and all its attached devices between partitions that share the same bus dynamically. The partition that currently owns the IOP must not be in use before you switch the IOP to another partition.

Switching requires removing the IOP from its current partition and adding it to a different one. In other words, two partitions cannot use the IOP and its devices at the same time.

Attention:

When switching disk unit IOPs, ensure that all disk units that belong to that specific IOP are first removed from the auxiliary storage pool and are in a non-configured status.

Logical partition concept: tagged resource

A tagged resource is an IOP that you select because it controls a device that performs a specific function for a logical partition. The devices that perform the essential functions are the alternate IPL device, partition console, Electronic Customer Support IOP, and load source resource.

Alternate IPL device

The media in the “Logical partition concept: alternate restart (IPL) device and removable media device” on page 13 is what the system uses to start from when you perform a D-source IPL. The device can be a tape drive or an optical device. The alternate IPL device loads the Licensed Internal Code contained on the removable media instead of the code on the load source.

Partition console

When using Operations Console, the “Logical partition concept: console” on page 14 and ECS IOP should be the same. The console is the first workstation that the system activates in the partition. The system assumes that this console will always be available for use.

Electronic Customer Support IOP

The Electronic Customer Support IOP is a communications IOP you can select to support Electronic Customer Support on a system or secondary partition. Electronic Customer Support is the part of the operating system that allows you to access the following:

- Question-and-answer (Q and A) function.
- Problem analysis, reporting, and management.
- Fixes (or Program temporary fixes, PTFs).
- IBM product information.
- Technical information exchange.

Load source resource

Each logical partition must have one disk unit designated as the “Logical partition concept: load source” on page 14. The load source resource is the IOP with the load source. The load source contains the Licensed Internal Code and the configuration data for the logical partitions. The system uses the load source to start the logical partition. The system always identifies this disk unit as unit number 1.

Logical partition concept: SPD and PCI

Hardware features are packaged in two different forms: System Product Division (SPD) or Peripheral Component Interface (PCI), depending upon the server model.

SPD I/O adapters (IOAs) are packaged with I/O processors (IOPs) and do not require a separate card slot. The device connects to the slot containing the IOA and the IOP.

PCI IOAs are packaged separately from the IOP and require a separate card slot. The IOP in a card slot connects to the IOA in another card slot. The device attaches to the IOA.

“Hardware for logical partitions” on page 4

Logical partition concept: Processor

A processor is a device that executes programmed instructions. Logical partitions support “Logical partition concept: dedicated processor” and “Logical partition concept: shared processor pool” on page 11. The more processors you have the greater number of concurrent operations being executed at any given time. Your processor sends and receives information from the different parts of the system (from hardware and software).

Processors can work as a group to decrease the computing time that an operation requires. The fewer the number of processors in a system the longer the amount of computing time needed. If more processors are allocated to a partition, the greater the number of concurrent operations.

You measure the total system performance in the Commercial Processing Workload (CPW) that is unique for each model. The relative performance of a partition equals the CPW for the entire system times the number of processors in a logical partition divided by the total number of processors on the system.

Relative logical partition performance = (CPW) (# processors in the logical partition/total # of processors).

From the Configure Logical Partition window, you can view all of the system processor hardware resources. From the primary partition, you can also see which processors the logical partitions own.

If a processor fails while the server is running, then all the logical partitions on that server (not just the one with the faulty processor) fail. If a processor failure is detected during the system restart (IPL), the logical partition configuration manager will attempt to honor the minimum processor setting for all partitions. Once minimums have been met, all remaining resources are distributed, proportional to the intended allocation, among the appropriate partitions. If the partition minimums cannot be met, all resources are left in the primary partition and no secondary partitions are started. An entry is placed in the primary partition Product Activity Log (PAL) with a System Reference Code (SRC) of B6005342 to indicate that the minimum configuration could not be met. The primary partition PAL also contains one or more entries to indicate the failing hardware. You can view processor errors in the Product Activity Log (PAL) on the primary partition.

“Hardware for logical partitions” on page 4


Logical partition concept: dedicated processor

Dedicated processors are whole processors that are dedicated to a single partition. The dedicated processor handles the processing for a specific logical partition.

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

To adjust to changing workloads, you can move dedicated processors within the minimum/maximum values you establish without needing to restart the partition. These values enable you to establish a range within which you can dynamically move the resource without needing to restart the logical partition. When

you change the minimum/maximum values, it requires you to restart the partition. Minimum values dictate what is required to restart the partition. If the minimum value is not met for all logical partitions, only the primary will restart.



For example, a server with 4 physical processors can have 3 logical partitions with two partitions having 1 dedicated processor and one partition having 2 dedicated processors.


“Hardware for logical partitions” on page 4

Logical partition concept: shared processor pool

The shared processing pool allows you to assign partial processors to a logical partition. The physical processors are held in the shared processing pool and are shared among the logical partitions. A minimum of 0.10 processing units can be configured for any partition using shared processors. The primary partition may require more than 0.10 processing units when starting secondary partitions, otherwise timeout conditions may occur with resources communicating directly with the primary partition. Each system processing capacity and partition configuration needs to be evaluated to determine reasonable processor units for primary and secondary partitions.

Virtual processors are the whole number of concurrent operations which the operating system can utilize. The processing power can be conceptualized as being spread equally across these virtual processors. Selecting the optimal number of virtual processors is dependent on the workload in the partition: some benefit from greater concurrence, others require greater power. It is recommended to keep a balance of virtual processors to processors units. If less than or equal to 1.00 processing units are specified, 1 virtual processor should be used. Likewise, if less than or equal to 2.00 processing units are specified, 2 virtual processors should be used. If an imbalance of processing units and virtual processors occurs, partition batch performance may be degraded.

To accommodate changing workloads, you can adjust shared processing units within the minimum/maximum values you establish without needing to restart the partition. These values enable you to establish a range within which you can dynamically move resources without needing to restart the logical partition. When you change the minimum/maximum values, you must restart the partition. Minimum values dictate what is required to restart the partition. If the minimum value is not met for all logical partitions, only the primary will restart.



For example, a system with 4 processors in the shared pool provides 4.00 processing units. Five logical partitions could distribute the processing power in the following way: Partition 0 has 2.00 processing units and 2 virtual processors, partition 1 has 0.50 processing units and 1 virtual processor, partition 2 has 0.50 processing units and 1 virtual processor, partition 3 has 0.75 processing units and 1 virtual processor, and partition 4 has 0.25 processing units and 1 virtual processor. The sum of the 5 logical partition's processing units is less than or equal to the total number of processing units in the shared pool. But the total number of virtual processors is 6.

“Hardware for logical partitions” on page 4

Logical partition concept: memory

“Logical partition concept: Processor” on page 10 use memory to temporarily hold information. Memory requirements for partitions are dependent on partition configuration, I/O resources assigned, and applications used. When you create a partition, you must add a whole number of megabytes of memory to it (1 MB = 1024 x 1024 bytes). The primary partition needs a minimum of 256 MB of memory. Depending on configuration values used, a primary partition may require more than 256 MB. Secondary partitions running V4R4 and V4R5 require a minimum value of 64MB. Secondary partitions running V5R1 and V5R2 require a minimum of 128 MB of memory. Depending on configuration values used, a secondary partition may require more than 128 MB.

Memory in each logical partition operates within its assigned minimum and maximum values. You can dynamically move memory between V5R1 and V5R2 logical partitions without needing to restart the affected partitions as long as your request for memory movement is within the range of the minimum and maximum values specified during partition creation. When you request to dynamically move memory between partitions, note that this memory is removed and added to the base memory pool of each of the partitions (*BASE pool). Private memory pools or shared memory pools are not affected. If the move request exceeds the amount of memory available in the base pool, the system will only free up excess memory pages after keeping the minimum amount of memory required in the base pool. This value is determined by the base storage minimum size (QBASPOOL) system value. To prevent any data loss during memory movement, the system first writes any data from memory pages to the disks before making the memory pages available to another partition. Depending on the amount of memory you have requested to move, this may take some time.

Each partition will report its runtime minimum memory size. This value is an estimate of how much memory in the partition is locked into that partition and cannot be moved dynamically. Reducing the number of processes or threads within a partition or changing the *BASE pool will affect the runtime minimum.

The full amount of memory that you assign to a logical partition may not be available for the partition's use. Static memory overhead required to support the assigned maximum memory will affect the reserved or hidden memory amount. This static memory overhead will also influence the minimum memory size of a partition.

The minimum memory size assigned to a partition can be changed with only a restart of the partition. However, changes to the maximum memory size require a restart of the entire system and may require a larger minimum memory value.

If there is a memory failure, the system will attempt to accommodate the minimums for all partitions. If all minimums are satisfied, the secondary partitions will restart with all additional resources distributed proportionally to their allocation. If all minimums are not satisfied, the system will put all resources in the primary and fail to restart with the partition configuration protected and a product activity log (PAL) error of B6005343. You can view memory errors in the PAL on the primary partition.

“Hardware for logical partitions” on page 4

Logical partition concept: disk units

Disk units store data. The server can use and reuse this data at any time. It is more permanent than memory, but you can still erase it.

You cannot separate the disk units on an “Logical partition concept: IOP” on page 8 into different logical partitions. You can create user auxiliary storage pools (ASPs) on any logical partition. However, you cannot create a cross-partition ASP. All of the disk units you assign to an ASP must be from the same logical partition. For more information on auxiliary storage pools (ASPs), refer to disk pools.

You can also create an independent auxiliary storage pool (ASP). An independent ASP is a collection of disk units that can be taken offline, or made unavailable, independent of other disk pools because the data in the independent ASP is self-contained. The independent ASP can also be brought online, or made available, while the system is active, without having to perform a restart. For more information on independent ASPs, refer to independent disk pools.

The server maintains the logical partition configuration data on the disk unit that you designate as the “Logical partition concept: load source” on page 14 for each logical partition.

When you move disk units, you may need to clear any logical partition configuration data that is contained on them.

“Hardware for logical partitions” on page 4

Logical partition concept: alternate restart (IPL) device and removable media device

A removable media device reads and writes to media (tape, CD-ROM, or DVD) . You must have either a tape or an optical (CD-ROM or DVD) device available for every logical partition. The system also uses some of these devices as the alternate restart or IPL device and alternate installation device.

Logical partitions can (depending on your hardware setup) share a tape or optical device and its attaching IOP among themselves. However, only one logical partition can use the device at any time. To switch devices between partitions, you must move the IOP, with the shared device, to the desired logical partition. For more information on how to move an IOP, refer to “Logical partition concept: IOP” on page 8.

Alternate IPL Device

The media in the device is what the system uses to start from when you perform a D source IPL. The alternate IPL device loads the Licensed Internal Code contained on the removable media instead of the code on the “Logical partition concept: load source”. It may also install the system.

“Hardware for logical partitions” on page 4

Logical partition concept: console

Each logical partition must have a console attached to it through an “Logical partition concept: IOP” on page 8. The console is the first workstation that the system activates. The system assumes that this console will always be available for use. You can access the dedicated service tools (DST) display only from this console.

A secondary partition console can be a twinaxial workstation, a local console on a network or a local console directly attached to the server.

Attention:

If you plan on using Operations Console local console on a network and you have a twinax IOA on the same IOP, it is possible the twinax workstation may come up first and become the console. Possible solutions would be to have the twinax IOA on a different IOP, to configure the terminal at a different address other than 0, or to unplug the device from the twinax cable.

For an Operations Console local console directly attached to the server configuration, the console IOP needs to be marked as both the console and electronic customer support (ECS) IOP.

For any other type of console, including Operations Console local console on a network configuration, Operations Console uses a token ring or Ethernet communication card, simply show the preferred type of IOP and select the desired one on the New Logical Partition - Console panel.

For information on Operations Console migration, refer to Plan for your Operations Console Migration.

“Hardware for logical partitions” on page 4

Logical partition concept: expansion unit

You can add an expansion unit to many of the iSeries servers to support additional features and devices. If you want to create logical partitions on your iSeries server, you will probably need to add an expansion unit. This will contain the additional hardware that you need for each logical partition.

There are different types of expansion units. Some expansion units can support only disk units (storage expansion unit), while others can support a variety of hardware (system expansion unit). This is dependent on the kind of buses and IOPs that are installed in the unit.

Expansion units generally contain one or two system I/O “Logical partition concept: bus” on page 5 with multiple “Logical partition concept: IOP” on page 8 that control various I/O devices.

“Hardware for logical partitions” on page 4

Logical partition concept: load source

Each logical partition must have one “Logical partition concept: disk units” on page 13 designated as the load source. The load source contains the Licensed Internal Code and the configuration data for the logical partitions. The server uses the load source to start the logical partition. The server always identifies this disk unit as unit number 1.

The logical partition configuration data on the load source of the primary partition is the master copy. The server uses this copy to check the integrity of the configuration data that it maintains on the load source of each logical partition.

Any time the load source of a logical partition is erased, the logical partition configuration data must be recovered. On a secondary partition, the server automatically rewrites the data by using the master copy from the primary partition. On the primary partition, you must manually recover the configuration data.

When you move the load source for a logical partition to a different server or different logical partition as a nonconfigured disk unit, you need to clear its configuration data. This recovery action corrects configuration data problems.

Logical partition concept: interactive performance

You can specify an interactive performance minimum value equal to the minimum amount of interactive performance needed to support the logical partition. The maximum value must be less than the amount of interactive performance available on the system.

Your server has a certain amount of interactive performance based on the type of server and the number of processors. Interactive performance refers to how much the user must interact with (respond to prompts from) the computer. You can contrast interactive with batch (no user intervention required).

Given the amount of interactive performance on your server (CPW), you need to determine what percentage will be available for each logical partition. The combined total of each logical partition's interactive performance must not exceed 100%.

When you create logical partitions, you specify what percentage of interactive performance each logical partition gets. The amount of interactive performance you can assign to a logical partition depends on the type of server and the number of processors in the partition. A logical partition may not be able to use the full range of the interactive performance. The server alerts you when you enter a value that is not correct.

You can set up interactive performance so that you can change the values between logical partitions without needing to restart the entire server using iSeries Navigator. The minimum and maximum specify the range of values you can enter without needing to restart the partition. When you change the limits for the logical partition's interactive performance, you must restart the partition.

Any unused portion of the interactive performance does not automatically transfer to other logical partitions.

Software licensing and licensed programs for logical partitions

Unique software resources exist and operate on hardware that is assigned to each partition. These software resources include separate copies of Licensed Internal Code, OS/400, and other licensed program products. Additionally, language feature codes, security, user data, most system values, and software release and fixes (or program temporary fixes, PTFs) remain unique for each logical partition.

Software license behavior varies by software product. Each solution provider has its own licensing strategy. You can use IBM software products, licensed by the processor group, in any partition. All you need to do is purchase one license for the iSeries server. You can then install the product on any partition you choose. IBM user-based products are priced based on the total number of users running across all partitions of an iSeries server.

Software licensing and pricing for IBM software products on servers running multiple partitions remains largely unchanged from current licensing and pricing policies. Processor-based products running in an environment with logical partitions are priced according to the Software Machine Group of the underlying iSeries hardware model. Processor-based one time charge IBM software products are licensed to run

simultaneously in all of the partitions on the server. IBM user-based products are priced based on the total number of users running across all partitions of an iSeries.

Software license management functions that currently exist in OS/400 are available in an environment with logical partitions. Software license management supports the various pricing models that iSeries Independent Software Vendors typically use to license their products on an iSeries server.

Software license management has 3 usage types: registered users, concurrent users and processors. All three count usage across the server. It can determine and enforce the number of users for a licensed product across multiple logical partitions of an iSeries.

Logical partition release support

Systems with logical partitions have the capability to support more than one version of OS/400. The logical partition strategy is to support up to three different releases on the same system. Using the primary partition as the reference release (denoted as P), the strategy is to support secondary partitions of one previous release of OS/400 (P - 1), the same release as the primary (P), and a future release (P + 1).

For example, if your primary partition is running V4R5, you can install V5R1 on a secondary partition. However, since your primary partition is running V4R5, you will not benefit from the V5R1 enhancements in the secondary partition running V5R1. You are limited to the release function of the primary partition. Likewise, if you install V5R1 on the primary partition, you can run V4R5 on one secondary partition while the other secondary partitions run V5R1. In this example, since the primary partition is running V5R1 all secondary partitions running V5R1 will have the increased functionality of the release. However, the secondary partition running V4R5 is limited to the release function established for V4R5.

OS/400 release support for the 6xx, 7xx, and Sx0 models

In general, 6xx, 7xx, and Sx0 hardware runs all software versions beginning with V4R4. Additionally, when this hardware runs V4R4 in the primary, V5R1 is also supported (P+2) in a secondary partition. These models must have 2 or more processors and are not capable of supporting the shared processor pool.

Primary	Secondary (P-1)	Secondary (P)	Secondary (P+1)	Secondary (P+2)
V4R4	V4R3 LPAR not supported	V4R4	V4R5	V5R1 exception for this release
V4R5	V4R4	V4R5	V5R1	not supported
V5R1	V4R5	V5R1	V5R2	not supported
V5R2	V5R1	V5R2	future release not supported	not supported

OS/400 release support for the 820, 830, 840, and 270 models

The 820, 830, and 840 hardware can support release V4R5 in a primary or secondary partition as long as the model has more than one processor. The 270 hardware can only support V4R5 in a secondary partition on the 2-way processor configuration. The 270 and 820, 830, and 840 models with a single processor can support V5R1 or later releases in all partitions.

Primary	Secondary (P-1)	Secondary (P)	Secondary (P+1)
V4R4 not supported	not supported	not supported	not supported

Primary	Secondary (P-1)	Secondary (P)	Secondary (P+1)
V4R5	V4R4 not supported	V4R5	V5R1
V5R1	V4R5	V5R1	V5R2
V5R2	V5R1	V5R2	future release supported

Attention:

Some 820, 830, and 840 models only support V5R1 or V5R2 in the primary partition. Consult your business partner, marketing representative, or service specialist for more details.

OS/400 release support for the 810, 825, 870, and 890 models

The 810, 825, 870, and 890 hardware can only support V5R2 in all partitions.

Primary	Secondary (P-1)	Secondary (P)	Secondary (P+1)
V5R2	not supported	V5R2	future release supported

OS/400 logical partition function by release

The capabilities available in OS/400 logical partitions vary with each operating system release. The primary partition release determines the base logical partition capability of the entire system. To utilize a specific capability, a secondary partition OS/400 release must also support the function. Use the following table to determine the logical partition function by OS/400 release.

Software Function	V4R4	V4R5	V5R1 and V5R2
Maximum Partitions	12 or the number of processors on the system, whichever is less.	12 for 6xx, 7xx, Sx0 models; 24 for 820, 830, and 840 models; or the number of processors on the system, whichever is less.	12 for 6xx, 7xx, Sx0 models; 10 times the number of system processors or 32 for 270, 8xx models. (The maximum number of partitions supported depends on the number of processors in the server model.)
Processors	<ul style="list-style-type: none"> Static: requires a restart of the partition to change. Dedicated to a partition. 	<ul style="list-style-type: none"> Static: requires a restart of the partition to change. Dedicated to a partition. 	<ul style="list-style-type: none"> Dynamic: may be changed with a restart of the partition. May be shared among multiple partitions.
Memory	Static: requires a restart of the partition to change.	Static: requires a restart of the partition to change.	Dynamic: may be changed without a restart of the partition.
Interactive	Static: requires a restart of the partition to change.	Static: requires a restart of the partition to change.	Dynamic: may be changed without a restart of the partition.


Software Function	V4R4	V4R5	V5R1 and V5R2
Virtual OptiConnect	<ul style="list-style-type: none"> • Static: requires a restart of the entire system to change. • Single network. 	<ul style="list-style-type: none"> • Static: requires a restart of the entire system to change. • Single network. 	<ul style="list-style-type: none"> • Dynamic: may be changed without a restart of the partition. • Single network.
Virtual Ethernet	Not supported.	Not supported.	<ul style="list-style-type: none"> • Dynamic: may be changed without a restart of the partition. • Up to 16 networks.
HSL OptiConnect	Not supported.	Not supported.	<ul style="list-style-type: none"> • Dynamic: may be changed without a restart of the partition. • May be shared among multiple partitions. • Single network.
I/O	<ul style="list-style-type: none"> • Allocated at bus-level or IOP-level. • IOP may be switched among partitions dynamically. • Bus ownership or bus usage (shared or dedicated) changes require a restart of the whole system. 	<ul style="list-style-type: none"> • Allocated at bus-level or IOP-level. • IOP may be switched among partitions dynamically. • Bus ownership or bus usage (shared or dedicated) changes require a restart of the whole system. 	<ul style="list-style-type: none"> • Allocated at bus-level or IOP-level. • IOP may be switched among partitions dynamically. • Bus ownership or bus usage (shared or dedicated) changes occur dynamically.
Guest Partition	Not supported.	Not supported.	Linux.

Once the software capabilities have been determined, ensure that all desired logical partition capabilities are supported by the particular hardware model. Additional details are located in Evaluate the hardware limitation of the iSeries .

Communication options for logical partitions


Logical partitions can use any of the following communication methods to interact with other partitions or servers.

Virtual Ethernet



Virtual Ethernet enables you to establish communication via TCP/IP between logical partitions. For each of the 16 ports enabled, the system will create a virtual Ethernet communications port, such as CMNxx with a resource type of 268C. Logical partitions assigned to the same virtual Ethernet then become available to communicate through that link. A physical system allows you to configure up to 16 different virtual local area networks. Virtual ethernet provides the same function as using a 1Gb Ethernet adapter. Token Ring or Ethernet 10Mbps and 100Mbps local area networks are not supported with virtual Ethernet. Virtual Ethernet requires V5R1 and later releases and can be used without any additional hardware or software.


High-Speed link (HSL) OptiConnect



High-Speed link (HSL) OptiConnect provides high speed system-to-system communication for PCI-based models. It requires standard HSL cables, but no additional hardware is required. To use HSL OptiConnect, you need to purchase OptiConnect for OS/400 software (a priced optional feature). The OptiConnect software will choose the Virtual OptiConnect path over a HSL or SPD OptiConnect external path if multiple paths are available.

You may enable HSL OptiConnect to other systems at any time for any partition within the system. But, you must install OptiConnect for OS/400 software before this feature can be used. When you enable or disable HSL OptiConnect, the changes take effect immediately.

Virtual OptiConnect




Virtual OptiConnect emulates external OptiConnect hardware by providing a virtual bus between logical partitions. You can use virtual OptiConnect without any additional hardware requirements. To use virtual OptiConnect, you only need to purchase OptiConnect for OS/400 (a priced optional feature).

The OptiConnect software will choose the virtual OptiConnect path over a HSL or SPD OptiConnect external path if multiple paths are available.

You may enable Virtual OptiConnect for a logical partition at any time. But, you must install OptiConnect for OS/400 software before this feature can be used. When you enable or disable Virtual OptiConnect, the changes take effect immediately.

SPD OptiConnect



SPD OptiConnect is a combination of hardware and software that allows you to connect multiple high-end iSeries servers through a high-speed, fiber-optic bus. This combination of optical bus speeds and efficient software makes OptiConnect a viable solution, providing multiple routes to the database. You can run both APPC or TCP/IP communications protocols over OptiConnect. TCP/IP support was added to OptiConnect beginning with OS/400 V4R4.

In order to use OptiConnect, you must purchase the hardware and OptiConnect for OS/400 (a priced feature).

For each logical partition that participates in external OptiConnect, you need a dedicated “Logical partition concept: bus” on page 5. You can not assign this bus as shared.

For more information about OptiConnect, see OptiConnect for OS/400  .

“Hardware for logical partitions” on page 4



Printed in U.S.A.