Microsoft

Windows Server 2008

# Dynamic Hardware Partitioning

White Paper

Published: January 2008

For the latest information, see www.microsoft.com.

# Contents

# Disclaimer

This is a preliminary document and may be changed substantially prior to final commercial release of the software described herein.

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This white paper is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS DOCUMENT.

# Abstract

Hardware support for dynamic hardware partitioning—the hot addition and replacement of processors and memory while a system is running—has been available on high-end server systems for several years. Partners such as Bull, Fujitsu, Fujitsu Siemens, Hewlett-Packard, Hitachi, IBM, NEC, and Unisys have requested support for dynamic hardware partitioning in the Windows Server operating system.

As the performance of mid-range server systems continues to increase with greater and greater numbers of cores per processor, these systems will soon be capable of running enterprise-class applications. To provide increased reliability, availability, and serviceability (RAS) when running business critical applications, Microsoft expects hardware support for dynamic hardware partitioning to become more commonly available on mid-range server systems in the future.

Therefore, Microsoft Windows Server® 2008 has been enhanced to support dynamic hardware partitioning features. Windows Server 2008 continues to support hot addition of memory, which was introduced in Windows Server 2003 Service Pack 1 (SP1). In addition, Windows Server 2008 supports the following features:

- Hot addition of processors
- Hot replacement of processors and memory
- Hot addition of I/O host bridges

Native support for PCI Express in Windows Server 2008 also enables the hot plug of PCI Express devices such as storage host bus adapters (HBAs) and network interface cards (NICs). That new support, together with Windows Server Multipath I/O (MPIO), provides further RAS benefits for mid-range server systems.

This paper provides information about dynamic hardware partitioning and the architecture that is supported by Windows Server 2008 and future versions of the Windows Server operating system.

## Audience

The primary audiences of this paper are IT professionals, systems administrators, and other decision makers who are considering the acquisition of new server systems that support dynamic hardware partitioning.

Technical information for independent hardware vendors (IHVs), independent software vendors (ISVs), and original equipment manufacturers (OEMs) is available on the Microsoft Web site. For more information, see the Resources section at the end of this white paper.

# Dynamic Hardware Partitioning Overview

## Concepts

A *hardware partitionable server* can be configured into one or more isolated *hardware partitions*. Each hardware partition in the server is assigned its own processors, memory, and I/O host bridges that are independent from all other hardware partitions in the server. Each hardware partition runs an independent instance of the operating system.

A hardware partition consists of one or more *partition units*. A partition unit is the smallest unit of hardware that can be assigned to a particular hardware partition. A partition unit can be an individual processor, memory module, or I/O host bridge, or it can be a hardware module or board that contains a combination of these components. Today's hardware-partitionable servers typically have multiple hardware components in each partition unit. For example, a single partition unit could consist of four processors, a memory module, and some I/O. In this situation, all of the hardware components in the partition unit must be added or replaced in a hardware partition as a single unit. As the number of processor cores per physical processor increases, the number of components per partition unit is likely to come down. However, with the memory controllers being implemented internal to the processors, processors and memory are likely to continue to be included in a single partition unit.

On a *statically partitionable server*, the configuration of partition units that are assigned to each hardware partition cannot be changed while the system is running. To change the configuration, the system must be powered down and restarted. Windows Server 2000 and later versions of Windows Server support statically partitionable servers.

On a *dynamically partitionable server*, the configuration of partition units that are assigned to a particular hardware partition can be changed while the system is running. Partition units can be added or replaced without restarting the instance of the operating system that is running on the hardware partition. Common operations include:

**hot add**
: Adding a partition unit to a running hardware partition.

**hot replace**
: Replacing a partition unit with an identical replacement partition unit that is already present in the system. Note that this is a single atomic operation that is not the same as a hot remove operation followed by a hot add operation.

**hot remove**
: Removing a partition unit from a running hardware partition.

Windows Server 2003 SP1 Enterprise Edition and Datacenter Edition support hot add of memory on x86-based, x64-based, and Itanium-based systems.

Windows Server 2008 supports hot add of processors, memory, and I/O host bridges plus hot replace of processors and memory on x64-based and Itanium-based systems.

**Note:** Although Windows Server 2008 does not support hot remove, Microsoft is considering including support for hot remove in a future version of Windows Server.

## Features

Windows Server 2008 supports the following dynamic hardware partitioning features:

- Processors, memory, and I/O host bridges can be hot added to a hardware partition while the system is running.
- Processors and memory can be hot replaced in a hardware partition while the system is running.
- Device drivers and applications can register to be notified of changes to the hardware configuration so that they can adjust their resource allocations accordingly.
- Existing applications should continue to run without modification. However, in order for an application to take advantage of any new hardware resources that are added to the hardware partition after the application has been started, the application must register for notification of changes to the hardware configuration.

The following table indicates which dynamic hardware partitioning features are available on each of the Windows Server 2008 editions.

| Feature<br><br>Windows Server 2008 Edition | Hot add memory | Hot add processor | Hot add I/O host bridge | Hot replace memory | Hot replace processor |
|---|---|---|---|---|---|
| **Standard x86** | | | X | | |
| **Enterprise x86** | X | | X | | |
| **Datacenter x86** | X | | X | | |
| **Standard x64** | | | X | | |
| **Enterprise x64** | X | | X | | |
| **Datacenter x64** | X | X | X | X | X |
| **For Itanium-Based Systems** | X | X | X | X | X |

For a specialty version of Windows Server 2008, such as Web Server, Storage Server, or Compute Server, the supported dynamic hardware partitioning features are based upon the base version of the operating system that is used for the specialty version.

Windows Server 2008 Datacenter and Windows Server 2008 For Itanium-Based Systems are licensed on a per-processor basis, not a per-instance basis. Therefore, you can configure a hardware partitionable server that is running either of these editions of Windows Server 2008 into as few or as many hardware partitions as the systems supports without requiring the purchasing of any additional operating system licenses. For more information, see the Operating System Licensing section of this white paper.

## Benefits

The following are benefits of dynamic hardware partitioning.

- You can add hardware resources (processors, memory, and I/O) to a running hardware partition without requiring the operating system to be shut down. This allows you to dynamically increase the amount of resources in the system when there is increased demand for those resources.
- You can replace hardware resources (processors and memory) in a running hardware partition without requiring the operating system to be shut down. This allows you to replace hardware that is failing without disrupting normal operation of the system.

The key advantage of a server that supports dynamic hardware partitioning is the increased system availability that results from being able to increase the amount of hardware resources or replace failing hardware resources in a running hardware partition without interrupting the operation of the server workload.

## Integration with the Windows Hardware Error Architecture

In versions of the Windows Server operating system prior to Windows Server 2008, the operating system supported several unrelated mechanisms for reporting hardware errors that were caused by the core platform hardware components (processors, memory, I/O, and so on). These mechanisms provided little support for error recovery. For uncorrected errors, the operating system simply stopped running and then recorded some of the available error information in the system event log after the system was restarted.

The ability to determine the root cause of platform hardware errors was hindered by the limited amount of error information that was recorded in the system event log. The operating system was not capable of preventing system failures caused by platform hardware errors because there was no common error record format or recovery mechanism and there was little support for hardware error management applications.

The Windows Hardware Error Architecture (WHEA), which was introduced in Windows Server 2008, extends the previous hardware error reporting mechanisms and brings them together as components of a coherent hardware error infrastructure.

WHEA takes advantage of the additional hardware error information that is available in today's hardware devices and integrates much more closely with the system firmware.

As a result, WHEA provides the following benefits:

- Allows for more extensive error data to be made available in a standard error record format for determining the root cause of hardware errors.
- Provides mechanisms to help recover from hardware errors to avoid causing the system to stop running when a hardware error is nonfatal or recoverable. This increases the overall RAS of the system.
- Supports user-mode error management applications and enables advanced system health monitoring through Event Tracing for Windows (ETW) events and by providing an API for error management and control.
- Provides extensibility, so that as hardware vendors add new and better hardware error reporting mechanisms to their devices, WHEA enables the operating system to gracefully accommodate the new mechanisms.

A system management application for a dynamically partitionable server can monitor the health of the system through the hardware error data that is provided by WHEA and automatically initiate the replacement of processors or memory modules that generate hardware errors. Integrating these two features of Windows Server 2008 ensures the continued stability and availability of the platform.

## Hardware Partitioning Compared to Server Virtualization

Server virtualization enables the creation of multiple virtual machines that share the same physical hardware. Each virtual machine runs a separate instance of an operating system. The hardware resources that are assigned to each virtual machine are managed by a virtual machine monitor (VMM) or hypervisor. By consolidating several low-utilization servers to all run in separate virtual machines in a single server, one can maximize the server utilization while reducing the physical space requirements for the server hardware. Consolidation also reduces equipment costs and lowers the electrical power consumption for server power and cooling.

Server virtualization products typically provide availability features through a clustering solution. Some virtualization solutions include a product feature, such as the Microsoft Quick Migration feature, that can save the complete state of the operating system in the virtual machine to disk and then restart the virtual machine on another system in just a few seconds. This transition is often fast enough that any established network connections between client systems and the applications or services that are running in the virtual machine do not lose connectivity.

Static hardware partitioning enables the creation of multiple isolated hardware partitions on a single server, each of which has its own dedicated hardware resources. Similar to server virtualization, each hardware partition runs a separate

instance of an operating system. However, the hardware resources that are assigned to each hardware partition are managed by a service processor, not by a VMM. A system management application (supplied by the hardware vendor) or a system console communicates with the service processor to enable a system administrator to set the hardware partition configuration. Where server virtualization is typically used to consolidate multiple lower-utilization workloads onto a commodity class server, hardware partitioning is typically used to consolidate multiple high-utilization (or scale-up) workloads onto a scale-up server.

If a hardware error occurs on a hardware partitionable server, it only affects the hardware partition that contains the hardware that caused the error because all of the hardware partitions in the server are isolated from each other.

Dynamic hardware partitioning adds the ability to change the configuration of the hardware resources that are assigned to each of the hardware partitions while the system is running. Dynamic hardware partitioning also increases the availability of the operating systems that are running on each of the hardware partitions because it supports hot replacement of physical processors and memory.

**Note:** You can combine server virtualization with hardware partitioning by installing a VMM in a hardware partition in a hardware partitionable server and creating multiple virtual machines within that hardware partition.

## Dynamic Hardware Partitioning Architecture

The following figure shows an example of a hardware partitionable server.

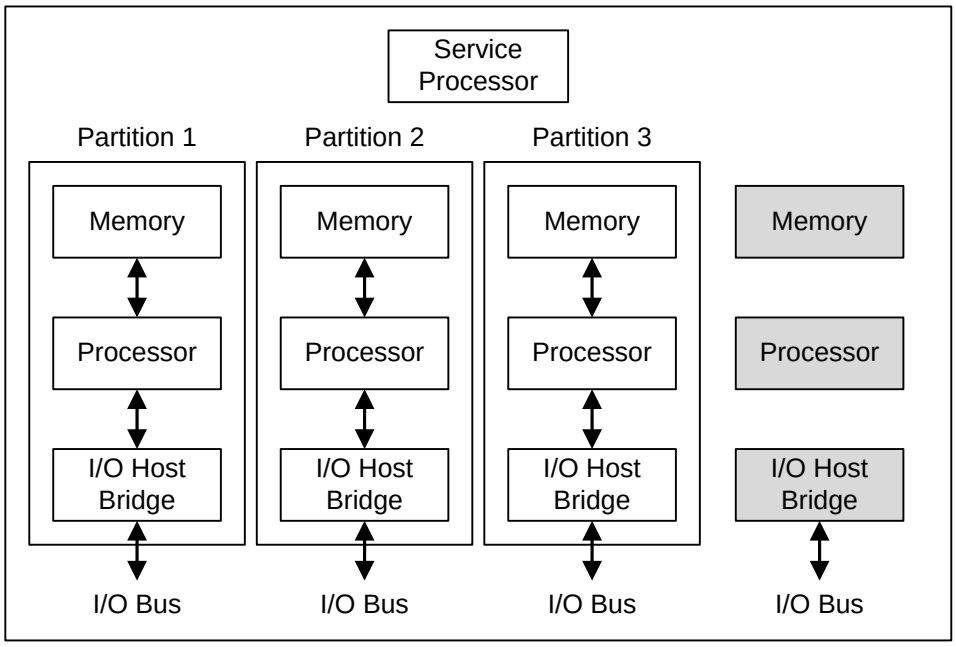


In this example, the server has a total of 12 partition units: 4 processors, 4 memory modules, and 4 I/O host bridges. One of each of these three types of partition units is assigned to one of three hardware partitions. The remaining partition units are available as spares that can be used to add hardware resources to one of the hardware partitions or to replace a failing partition unit in one of the hardware partitions.

In this example each partition unit consists of a single hardware component (a processor, memory module, or I/O host bridge). However, today's hardware partitionable servers typically have multiple hardware components in each partition unit. This simplified example is used for illustration purposes.

Each hardware partition is totally isolated from the other hardware partitions and runs an independent instance of the operating system. The service processor is responsible for managing the configuration of the hardware partitions. It controls the mapping of each of the partition units to a particular hardware partition and creates isolation between the different hardware partitions. In a dynamically partitionable server, this mapping can be reconfigured while the server is running. A system management application (supplied by the hardware vendor) communicates with the service processor to set the hardware partition configuration.

In Windows Server 2008, processors and memory are modeled as Plug and Play devices. This enables processor and memory devices to be started after the operating system has booted. This also enables device drivers and applications that

are running on Windows Server 2008 to register for Plug and Play notifications so that they can be notified when processors or memory have been added. For example, Microsoft SQL Server™ 2008 supports dynamic hardware partitioning.

I/O host bridges have been modeled as Plug and Play devices since Windows 2000.

## Dynamic Hardware Partitioning Scenarios

The following describes each of the hot add and hot replace scenarios that are supported by Windows Server 2008. Each of these scenarios begins with a dynamically partitionable server that is configured according to the following figure.

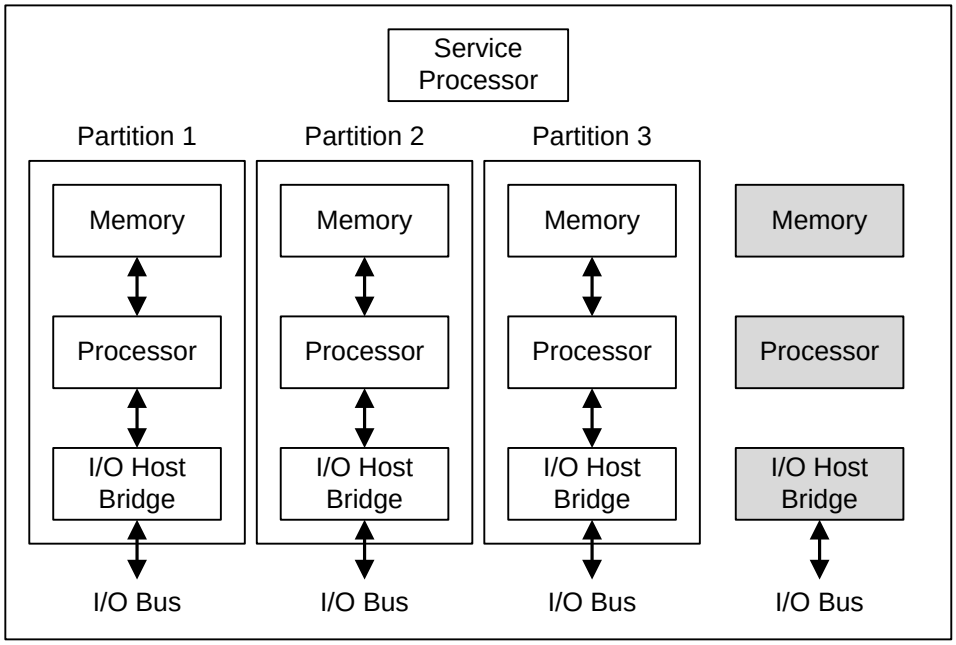


In this example each partition unit consists of a single hardware component (a processor, memory module, or I/O host bridge). However, today's hardware partitionable servers typically have multiple hardware components in each partition unit. This simplified example is used here for illustration purposes.

### Hot Add

When a hot add operation is performed on a dynamically partitionable server, a new partition unit is added to a running hardware partition. Windows Server 2008 supports hot add operations for processors, memory, and I/O host bridges.

Additional hardware resources are typically added to a running hardware partition when the utilization of the existing resources exceeds a level defined by the administrator of the server. For example, a new processor might be added to a hardware partition if the average processor utilization exceeds a specified threshold. Similarly, additional memory might be added to a hardware partition if the available physical memory goes below a specified threshold.

The procedure for changing the configuration of the hardware partitions in a dynamically partitionable server might be automatic in response to detecting a need for additional resources or might require intervention by a system administrator,

depending on the system management software that is provided by the hardware vendor.

For the purpose of describing what happens during a hot add operation for each type of hardware component, the following scenarios assume that the partition unit that is added to the hardware partition contains a single hardware component. In a typical case where a partition unit contains multiple hardware components, all of the hardware components that are contained in the partition unit are added to the hardware partition at the same time. In that situation, the scenario is really a combination of the steps that are described in these example scenarios.

In each of the following hot add scenarios, one of the spare partition units is added to partition number 3. However, the spare partition unit could have been added to any of the three hardware partitions in the server.

**Hot Add of a Processor**

When a condition that warrants adding another processor to a running hardware partition exists, a system management application (supplied by the hardware vendor) initiates the addition of an available spare processor to the hardware partition either automatically or with administrator intervention. It does this by communicating with the service processor, which in turn performs the following tasks:

1. Selects an available spare processor to add to the hardware partition.
2. Powers up and initializes the selected spare processor.
3. Adds the selected spare processor to the desired hardware partition.
4. Notifies Windows Server 2008 that a new processor has been added.

When Windows Server 2008 is notified that a new processor has been added to the hardware partition, it starts the new processor and adds it to the pool of available processors for executing threads. It also notifies any registered device drivers or applications that a new processor has been added to the hardware partition so that they can adjust their resource allocation for safe and optimal operation.

As well as notifying registered device drivers and applications, Windows Server 2008 also initiates a system-wide resource rebalance after adding a new processor to a hardware partition. This allows any device drivers that participate in the resource rebalance to disconnect and reconnect their interrupt handlers so that they can start using the new processor for handling hardware interrupts that are generated by the devices that the device drivers service. However, if participating in such a resource rebalance would cause a disruption in the operation of the applications or services that are running on the server, a device driver can be configured so that it does not participate in the resource rebalance.

For example, by default all of the device drivers for network interface cards (NICs) do not participate in such a resource rebalance so that all of the network connections that were established with the server prior to the hot add operation remain intact. For more information about how to configure the policy for a device driver that

determines whether the device driver participates in resource rebalancing following the dynamic addition of a new processor, see the documentation for the DEVPKEY_Device_DHP_Rebalance_Policy device property in the Microsoft Windows Driver Kit (WDK).

After the hot add operation is complete, Windows Server 2008 notifies the service processor, which in turn notifies the system management application. Portions of this process might be performed by the system firmware, depending on the design and architecture of the server.

After the hot add operation is complete, the example server system is configured as shown in the following figure, where the spare processor has been added to partition number 3.

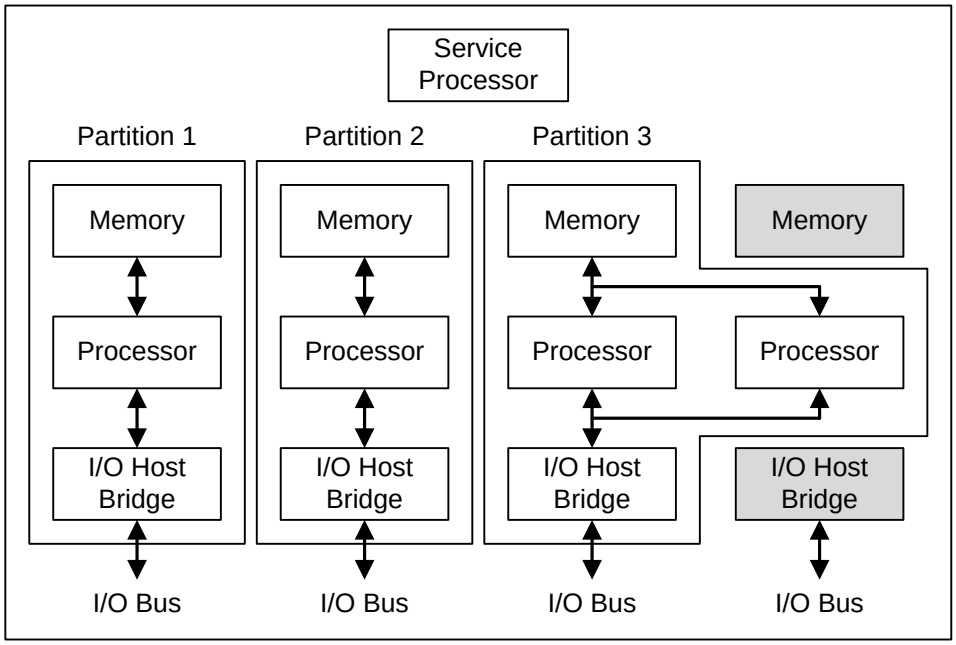


**Hot Add of Memory**

When a condition that warrants adding more memory to a running hardware partition exists, a system management application (supplied by the hardware vendor) initiates the addition of an available spare memory module to the hardware partition either automatically or with administrator intervention. It does this by communicating with the service processor, which in turn performs the following tasks:

1. Selects an available spare memory module to add to the hardware partition.
2. Powers up and initializes the selected spare memory module.
3. Adds the selected spare memory module to the desired hardware partition.
4. Notifies Windows Server 2008 that additional memory has been added.

When Windows Server 2008 has been notified that additional memory has been added to the hardware partition, it initializes the memory and adds it to the pool of available physical memory for applications to use. It also notifies any registered

device drivers or applications that additional memory has been added to the hardware partition so that they can adjust their resource allocation for safe and optimal operation.

After the hot add operation is complete, Windows Server 2008 notifies the service processor, which in turn notifies the system management application. Portions of this process might be performed by the system firmware, depending on the design and architecture of the server.

After the hot add operation is complete, the example server system is configured as shown in the following figure, where the spare memory has been added to partition number 3.

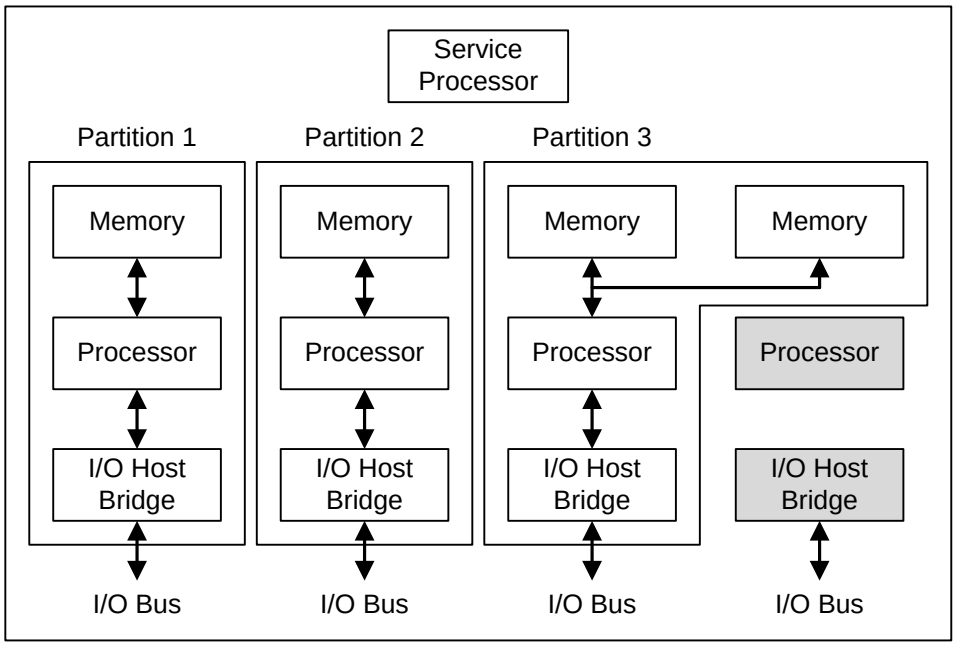


**Hot Add of an I/O Host Bridge**

When a condition that warrants adding another I/O host bridge to a running hardware partition exists, a system management application (supplied by the hardware vendor) initiates the addition of an available spare I/O host bridge to the hardware partition either automatically or with administrator intervention. It does this by communicating with the service processor, which in turn performs the following tasks:

1. Selects an available spare I/O host bridge to add to the desired hardware partition.
2. Powers up and initializes the selected spare I/O host bridge.
3. Adds the selected spare I/O host bridge to the desired hardware partition.
4. Notifies Windows Server 2008 that a new I/O host bridge has been added.

When Windows Server 2008 has been notified that a new I/O host bridge has been added to the hardware partition, it initializes the bridge and scans for any devices that are attached to the new I/O bus. For each device that is found on the bus, the

operating system loads the appropriate device drivers per the normal Plug and Play mechanism. Windows Server 2008 might also initiate a resource rebalance after a new I/O host bridge has been added to a hardware partition if resource arbitration for the new I/O host bridge or for any of the devices that are attached to the new I/O bus does not initially succeed.

After the hot add operation is complete, Windows Server 2008 notifies the service processor, which in turn notifies the system management application. Portions of this process might be performed by the system firmware, depending on the design and architecture of the server.

After the hot add operation is complete, the example server system is configured as shown in the following figure, where the spare I/O host bridge has been added to partition number 3.

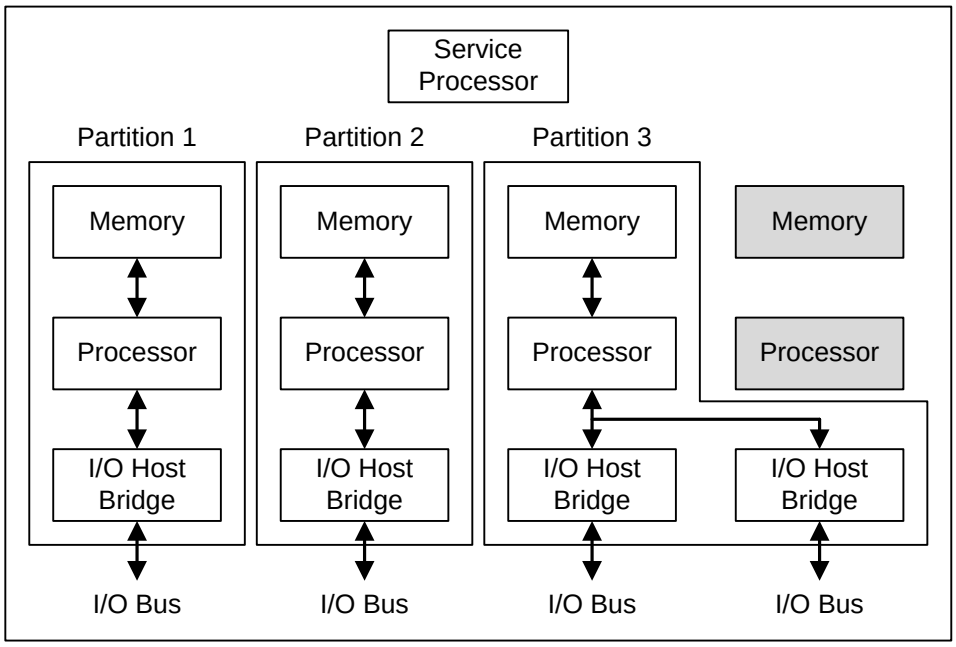


## Hot Replace

When a hot replace operation is performed on a dynamically partitionable server, an existing partition unit in a running hardware partition is replaced with an identical spare partition unit. Windows Server 2008 supports hot replace operations for processors and memory. Hot replace of I/O host bridges is not supported. When performing a hot replace operation, the replacement partition unit must be identical to the partition unit that it replaces.

A partition unit is typically replaced in a running hardware partition when an imminent failure of the partition unit is detected. For example, if the Windows Hardware Error Architecture reports more than a certain number of nonfatal and recoverable errors

that are caused by a particular processor or memory module, that partition unit can be scheduled for replacement before the system fails due to an unrecoverable error.

The procedure for changing the configuration of the hardware partitions in a dynamically partitionable server might be automatic in response to detecting a certain number of hardware errors or might require intervention by an administrator, depending on the system management software that is provided by the hardware vendor.

During a hot replace operation, the operating system temporarily puts the system into a *pseudo*-S4 sleep state. A pseudo-S4 sleep state is the same as a normal S4 (hibernate) sleep state except that the operating system does not save a hibernation file to the hard disk or power down the system. By putting the system into this special sleep state, the operating system stops all processing and I/O operations while the partition unit is replaced.

For the purpose of describing what happens during a hot replace operation for each type of hardware component, the following scenarios assume that the partition unit that is replaced in the hardware partition contains a single hardware component. In a typical case where a partition unit contains multiple hardware components, all of the hardware components that are contained in the partition unit are replaced in the hardware partition at the same time. In that situation, the scenario is really a combination of the steps that are described in these example scenarios.

In each of the following scenarios, one of the spare partition units replaces an existing partition unit in partition number 3. However, that the spare partition unit could have replaced a partition unit in any of the three hardware partitions in the server.

**Hot Replace of a Processor**

When a condition that warrants replacing a processor in a running hardware partition exists, a system management application (supplied by the hardware vendor) initiates the replacement of the processor with an available spare processor either automatically or with administrator intervention. It does this by communicating with the service processor, which in turn performs the following tasks:

1. Selects an available spare processor as the replacement processor.
2. Powers up and initializes the selected spare processor.
3. Notifies Windows Server 2008 that a replace operation is about to occur. In response to this notification, the operating system puts the system into a pseudo-S4 sleep state.
4. In cooperation with the operating system and the system firmware, copies the processor state from the failing processor to the replacement processor.
5. Maps the replacement processor into the hardware partition in place of the failing processor and removes the failing processor from the hardware partition.

6. Notifies Windows Server 2008 that the replace operation is complete. In response to this notification, the operating system wakes up from the pseudo-S4 sleep state and resumes normal processing.
7. Powers down the failing processor.
8. Notifies the system management application that the replace operation is complete.

Portions of this process might be performed by the system firmware, depending on the design and architecture of the server.

When a processor is hot replaced in a hardware partition, the hot replace operation is transparent to all of the device drivers and applications that are currently running when the replacement occurs, other than temporarily putting the system into the pseudo-S4 sleep state. Depending on the server architecture and the network timeout settings, the replace operation might occur fast enough that any established network connections between client systems and the applications or services that are running in the hardware partition do not time out.

After the hot replace operation is complete, the example server system is configured as shown in the following figure, where the original processor that was in partition number 3 is replaced with the spare processor.

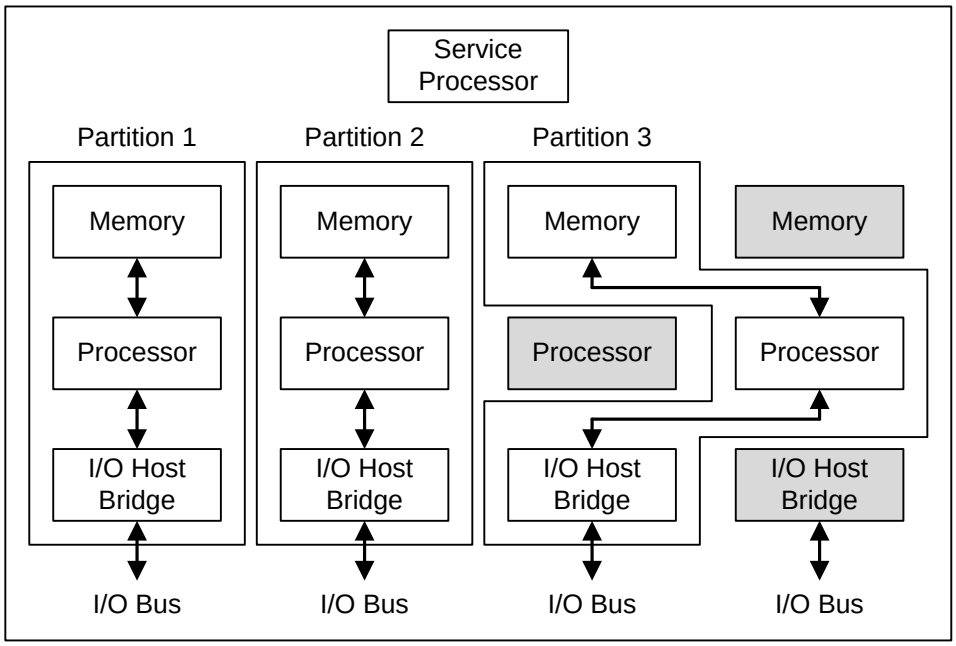

**Hot Replace of Memory**

When a condition that warrants replacing a memory module in a running hardware partition exists, a system management application (supplied by the hardware vendor) initiates the replacement of the memory module with an available spare memory module either automatically or with administrator intervention. It does this by communicating with the service processor, which in turn performs the following tasks:

1. Selects an available spare memory module as the replacement memory module.
2. Powers up and initializes the selected spare memory module.
3. Copies the memory contents from the failing memory module to the replacement memory module. Depending on the server architecture, this copy operation might be done with hardware assistance.
4. Notifies Windows Server 2008 that a replace operation is about to occur. In response to this notification, the operating system puts the system into a pseudo-S4 sleep state.
5. Copies the memory contents from the failing memory module to the replacement memory module for any memory locations that were subsequently changed after the copy operation in step 3 was initiated. Depending on the server architecture, the contents of the replacement memory module might stay in sync with the contents of the failing memory module, in which case this step is not necessary.
6. Maps the replacement memory module into the hardware partition in place of the failing memory module and removes the failing memory module from the hardware partition.
7. Notifies Windows Server 2008 that the replace operation is complete. In response to this notification, the operating system wakes up from the pseudo-S4 sleep state and resumes normal processing.
8. Powers down the failing memory module.
9. Notifies the system management application that the replace operation is complete.

Portions of this process might be performed by the system firmware, depending on the design and architecture of the server.

When a memory module is hot replaced in a hardware partition, the hot replace operation is transparent to all of the device drivers and applications that are currently running when the replacement occurs, other than temporarily putting the system into the pseudo-S4 sleep state. Depending upon the server architecture and the network timeout settings, the replace operation might occur fast enough that any established network connections between client systems and the applications or services that are running in the hardware partition do not time out.

After the hot replace operation is complete, the example server system is configured as shown in the following figure, where the original memory module that was in partition number 3 is replaced with the spare memory module.

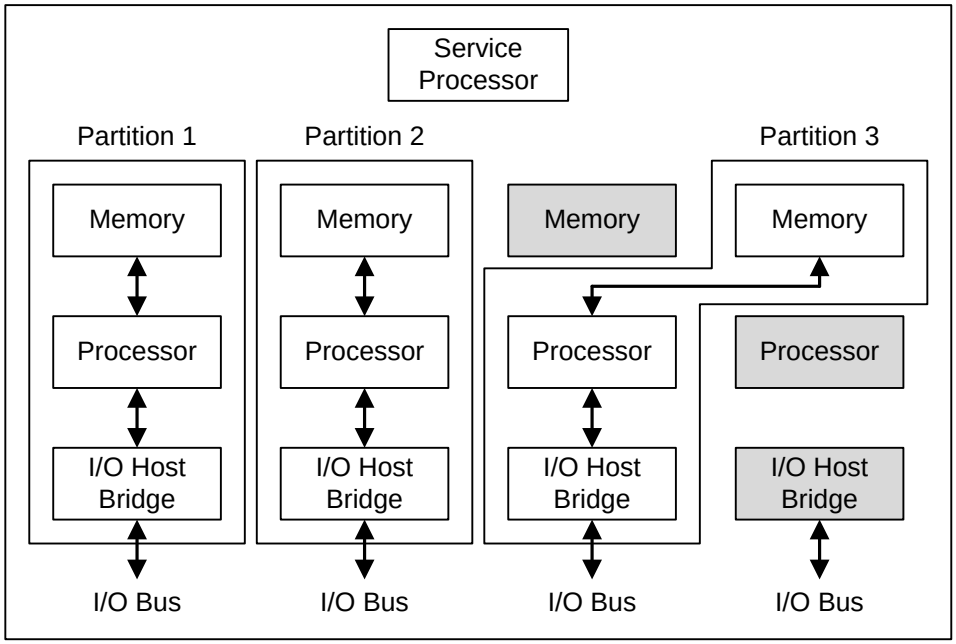

# Need for Device Drivers to Support Dynamic Hardware Partitioning

Certain assumptions that might have been made in the past by device driver developers are not valid on dynamically partitionable servers. Device drivers that are not designed with dynamic hardware partitioning in mind can cause data corruption, a system STOP error, or the system to hang if they are run on a dynamically partitionable server that executes a hot add or hot replace operation.

Therefore, customers should only use device drivers that have been specifically tested for Windows Server 2008 to ensure that there are no compatibility issues when run on a dynamically partitionable server.

**Recommendations**

Microsoft recommends as a best practice that vendors that develop device drivers observe the critical issues that are identified in the *Driver Compatibility for Dynamic Hardware Partitioning* white paper, even if a device driver is not specifically designed to be run on server systems. The Resources section at the end of this white paper includes a URL where the *Driver Compatibility for Dynamic Hardware Partitioning* white paper can be downloaded from the Microsoft Web site.

Although Windows Server 2008 does not support hot remove, Microsoft recommends that vendors who develop device drivers consider driver design issues that are related to hot remove when updating their device drivers to support hot add and hot replace. Microsoft is considering including support for hot remove in a future version of Windows Server.

For more information about how to add support for dynamic hardware partitioning to device drivers and applications, see the Dynamic Hardware Partitioning section in the Microsoft Windows Driver Kit (WDK) documentation. The Resources section at the end of this white paper includes a URL where the Dynamic Hardware Partitioning section of the WDK documentation can be accessed on the Microsoft Developer Network (MSDN) Web site.

# Purchasing a System that Supports Dynamic Hardware Partitioning


Every server system must pass the Windows Logo Kit (WLK) Server Logo tests in order to receive a **Certified for Windows Server 2008** Logo. Dynamically partitionable servers can receive an additional qualification for dynamic hardware partitioning support if they also pass the WLK Dynamic Hardware Partitioning tests.

Only those server systems that have received a Windows Server Logo are listed in the Windows Server Catalog and are supported by Microsoft. Dynamically partitionable servers that have passed the WLK Dynamic Hardware Partitioning tests are noted as such in the Windows Server Catalog to differentiate them from servers that do not support dynamic hardware partitioning.

For more information about the Windows Logo Program, see the Resources section at the end of this white paper.

For a list of servers that have received a Certified for Windows Server 2008 Logo and that support dynamic hardware partitioning, see the Windows Server Catalog online at www.windowsservercatalog.com.

# Product Licensing

## Operating System Licensing

Windows Server 2008 Standard and Windows Server 2008 Enterprise are licensed on a per-instance basis. Therefore, to run these editions of the operating system on a hardware partitionable server, an operating system license is required for each hardware partition.

Windows Server 2008 Datacenter and Windows Server 2008 For Itanium-Based Systems are licensed on a per-processor basis, not a per-instance basis. Therefore, for these editions of the operating system you need an operating system license for the total number of processors in the server. You can configure a hardware partitionable server that is running either of these editions of the operating system into as few or as many hardware partitions as the system supports without requiring the purchase of any additional operating system licenses.

**Note**: Each physical processor device can have multiple processor cores internally. However, the licensing for Windows Server 2008 Datacenter and Windows Server 2008 For Itanium-Based Systems is based on the number of physical processor devices in the system, not the number of processor cores.

For specialty versions of Windows Server 2008, such as Web Server, Storage Server, or Compute Server, the operating system licensing is based upon the base version of the operating system that is used for the specialty version.

## Application Licensing

### Microsoft SQL Server

SQL Server is licensed either on per-instance (server license plus client access licenses) basis, or on a per-processor basis. For the case of per-instance licensing, SQL Server has no limitations regarding the number of processors in the system. Therefore, when SQL Server runs in a hardware partition on a dynamically partitionable server, additional processors can be added to the hardware partition without violating the SQL Server license agreement.

For the case of per-processor licensing, SQL Server is only authorized to run in a system that contains no more processors than the number of SQL Server processor licenses. In this situation, when SQL Server runs in a hardware partition on a dynamically partitionable server, additional processors can only be added to the hardware partition until the number of processors equals the number of purchased SQL Server processor licenses without violating the SQL Server license agreement. Therefore, a customer should purchase SQL Server processor licenses for the maximum number of processors that would ever be added to the hardware partition,

which could mean purchasing a SQL Server processor license for every processor in the system.

A customer can choose which SQL Server licensing model they would like to use when they purchase the SQL Server product.

**Other Microsoft Server Applications**

Most other Microsoft server application products are licensed on a per-instance basis, where there are no limitations regarding the number of processors in the system. Therefore, when these products are run in a hardware partition on a dynamically partitionable server, additional processors can be added to the hardware partition without violating the product's license agreement.

Microsoft can only provide licensing information for Microsoft products. Consult your application software vendor for licensing information for any third-party products.

## Summary

Windows Server 2008 supports hot add and hot replace of processors and memory on dynamically partitionable server systems. This support for dynamic hardware partitioning provides an increased level of reliability, availability, and serviceability (RAS) as well as increased manageability and flexibility for these systems, resulting in greater uptime for business-critical applications.

## Resources

Dynamic Hardware Partitioning

http://www.microsoft.com/whdc/system/platform/server/dhp.mspx

http://msdn2.microsoft.com/en-us/library/bb742870.aspx

Driver Compatibility for Dynamic Hardware Partitioning

http://www.microsoft.com/whdc/system/platform/server/dhpbrief.mspx

Windows Hardware Error Architecture

http://www.microsoft.com/whdc/system/pnppwr/WHEA/default.mspx

http://msdn2.microsoft.com/en-us/library/aa939043.aspx

Virtualization

http://www.microsoft.com/virtualization/default.mspx

http://www.microsoft.com/whdc/system/platform/virtual/default.mspx

Windows Logo Program

http://www.microsoft.com/whdc/winlogo/default.mspx

Windows Server Catalog

http://www.windowsservercatalog.com

# Glossary

The following are definitions of terms related to dynamic hardware partitioning that are used in this white paper.

**partition unit**
The smallest unit of hardware that can be assigned to a hardware partition. This can be an individual processor, memory module, or I/O host bridge, or it can be a hardware module or board that contains a combination of these components.

**hardware partition**
A collection of partition units that is isolated from all other partition units in a system.

**hardware partitionable server**
A server that can be configured into one or more isolated hardware partitions.

**statically partitionable server**
A hardware partitionable server where the configuration of the partition units that are assigned to each hardware partition can be changed only when the system is shut down.

**dynamically partitionable server**
A hardware partitionable server where the configuration of the partition units that are assigned to each hardware partition can be changed while the system is running.

**hot add**
Adding a partition unit to a running hardware partition.

**hot replace**
Replacing a partition unit with an identical replacement partition unit in a running hardware partition. This is an atomic operation that is not the same as a hot remove operation followed by a hot add operation.

**hot remove**
Removing a partition unit from a running hardware partition.