
Contents

Preface	v	Chapter 3. WebSphere Virtual Enterprise in a micro-partitioned environment	9
Chapter 1. Virtualization and WebSphere Virtual Enterprise.	1	Configuring AIX 5.3 and AIX 6.1 on POWER5 and POWER6 (Micro-Partitioning)	11
Chapter 2. Supported server virtualization environments	5	Notices	13

Preface

This version of the white paper is for WebSphere® Virtual Enterprise *Version 6.1.1*.

If you are using Version 6.1.0.5, you can download the Version 6.1 white paper from the WebSphere Virtual Enterprise wiki.

This information changes frequently. Download the latest copy of this document on the WebSphere Virtual Enterprise Wiki or view the information in the WebSphere Extended Deployment information center.

Chapter 1. Virtualization and WebSphere Virtual Enterprise

By configuring application infrastructure virtualization in WebSphere Virtual Enterprise, you can pool together resources that are normally kept separate to accommodate the fluctuations of workload in your environment and increase the quality of service. You can also use application infrastructure virtualization with server virtualization capabilities that are provided by the physical hardware on which WebSphere Virtual Enterprise is hosted.

Application infrastructure virtualization

With *application infrastructure virtualization*, you can separate applications from the physical infrastructure on which they are hosted. Workloads can then be dynamically placed and migrated across a pool of application server resources, which allows the infrastructure to dynamically adapt and respond to business needs. Requests are prioritized and intelligently routed to respond to the most critical applications and users.

Typically, applications and Java™ 2 Platform, Enterprise Edition (J2EE) resources are statically bound to a specific server. Some of these applications might experience periodic increases in load that last a short time. The most costly time for an application to become unavailable is during a period of high demand. You must build your IT infrastructures to be able to accommodate these peaks. During the majority of time when your systems experience normal load, a large percentage of your computing capacity might go unused, making inefficient use of IT investments.

In a static environment, applications often span multiple enterprise archives (.ear files), and are not comprehensively defined so that the application can be portable between environments. Statically deployed applications rely on information that is found in the server to which they are deployed.

In the virtualized dynamic operations environment of WebSphere Virtual Enterprise, the static relationship is replaced with a dynamic relationship with looser coupling of applications or resources and server instances. Instead of statically binding applications to servers or clusters, you deploy applications to dynamic clusters, which are application deployment targets that can expand and contract depending on the workload in your environment.

After you deploy your applications to be mobile by using dynamic clusters, the placement of the applications is determined by the operational policies that you define. Autonomic managers control the placement of the server instances and how workload is routed for each application. If workload increases for a specific application, the number of server instances for the dynamic cluster that is hosting the application can increase, using available resources from other applications that are not experiencing increased workload.

Application infrastructure virtualization benefits:

- **Improved management of software and applications:** Management processes become more repeatable and less error-prone by using automated services and operational policies.

- **Allocation of software resources:** Dynamic reallocation of resources can occur based on shifting distributions of load among applications.
- **Increased number of applications:** More applications can run in a virtualized application environment than in a static configuration.
- **Reduced configuration complexity:** Loosened coupling between applications and the application server instances reduces the overall complexity and provides for a better, more usable environment.

Application infrastructure virtualization example




Figure 1. Application infrastructure virtualization in a WebSphere Virtual Enterprise environment. You deploy an application to a dynamic cluster that has a specified membership policy or node group. You do not deploy your applications to specific application servers. Instead, the application placement controller (APC) starts application server instances for your dynamic cluster based on the settings that you chose for the dynamic cluster.

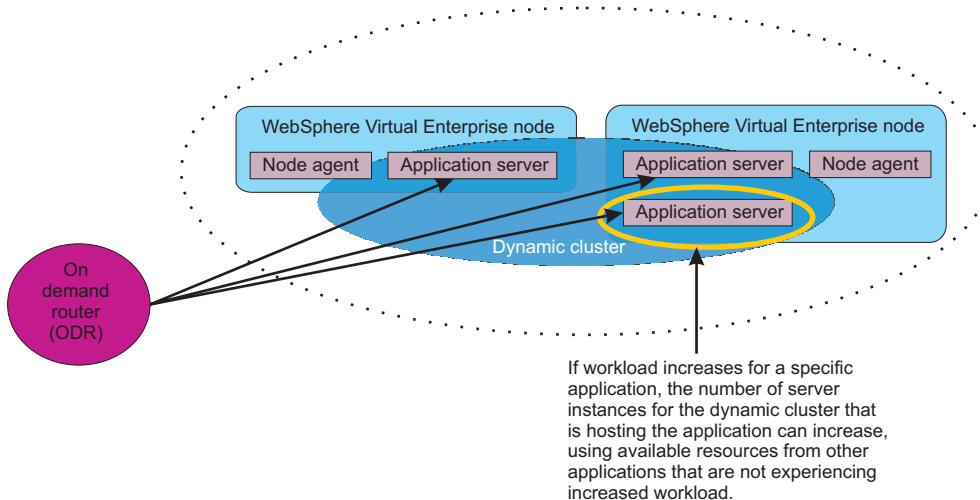


Figure 2. The starting of an additional application server to react to changes in application load. Additional application servers can start on the nodes that are selected by your dynamic cluster membership policy to handle additional requests that are coming in for the application.

Server virtualization

While WebSphere Virtual Enterprise provides virtualization of applications in your environment, you can also deploy WebSphere Virtual Enterprise on virtualized hardware, such as ESX servers, to take advantage of the capabilities provided by the hosting environment.

Server virtualization benefits:

- **Reduced amount of hardware in your environment:** You can run multiple WebSphere Virtual Enterprise nodes on the same physical hardware.

- **Improved hardware management:** You can more easily manage your environment because you have fewer physical computers and can use the server virtualization software to manage your images.
- **High availability of hardware:** By configuring server failover, your physical hardware can be highly available. When one server fails, it can be replaced by another server.
- **Dynamic allocation of hardware:** The physical resources, such as processors and memory, on your hosting computers can be shared among the virtual servers in your environment and dynamically allocated as needed. Because the resources are dynamically allocated, restarting the servers is not necessary.
- **Shared storage:** Multiple virtual servers or logical partitions can share the same physical storage. You do not need a physical hard drive for each virtual machine or LPAR.

WebSphere Virtual Enterprise in an environment with server virtualization

WebSphere Virtual Enterprise can operate in supported virtualized server environments. Different server vendors provide different virtualization capabilities, so the behavior of WebSphere Virtual Enterprise in different server virtualization environments can vary. However, common themes exist in server virtualization environments, such as the ability to share server resources across the virtual servers or logical partitions. Server virtualization environments can run in *shared processor mode* or *dedicated processor mode*. When you use shared processor mode, the physical processors are pooled and shared between the servers or logical partitions that are running on the physical computer. When you use dedicated processor mode, the physical processors are statically assigned to each virtual server or logical partition.

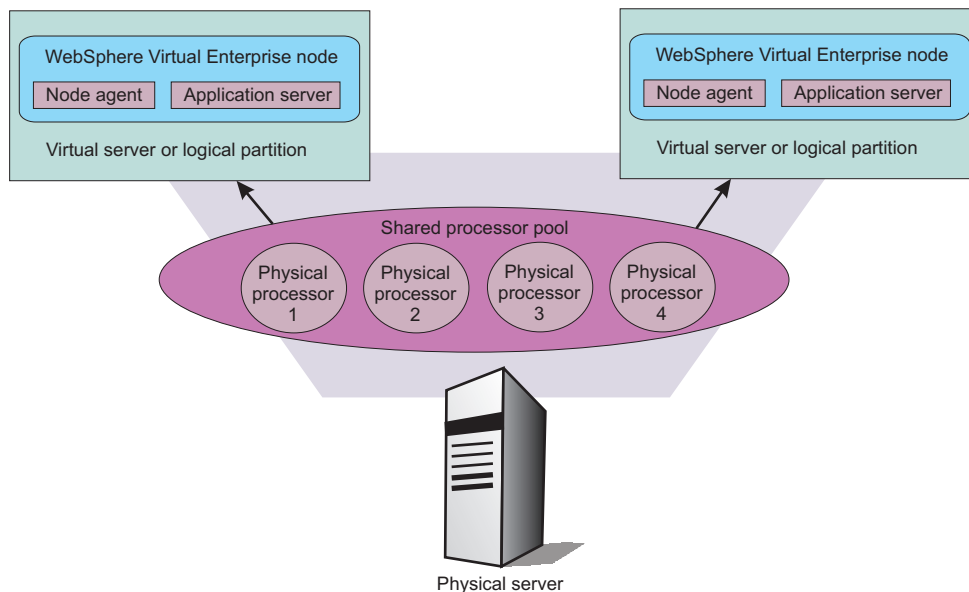


Figure 3. *Shared processor mode*. In shared processor mode, the physical processors are pooled and shared among the virtual servers or logical partitions.

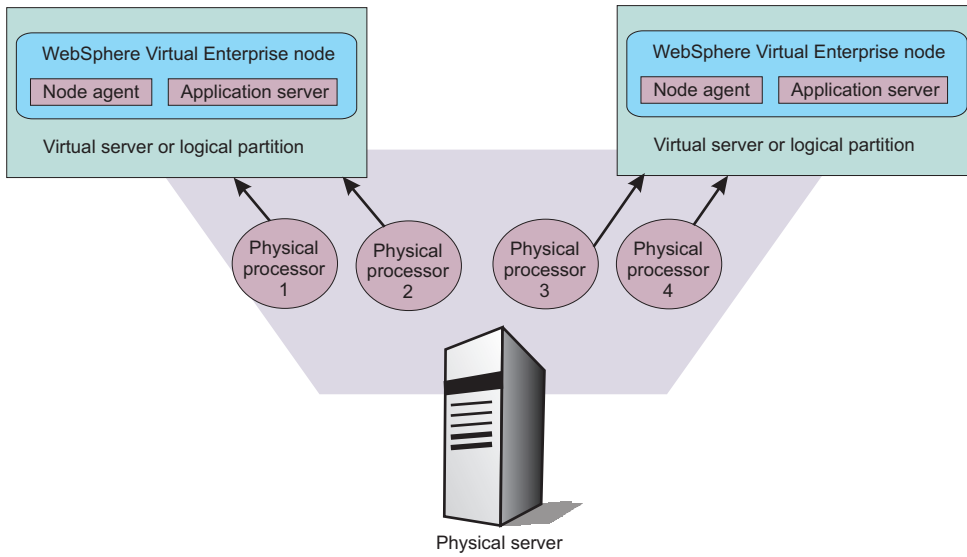


Figure 4. Dedicated processor mode. In dedicated processor mode, the physical processors are statically assigned to each virtual server or logical partition.

WebSphere Virtual Enterprise can also run in server virtualization environments with dedicated processor mode. The processor capacity is statically fixed to each virtual server or logical partition. The capacity and assignment do not change dynamically. Because the processor resource does not change for each virtual server or logical partition, using dedicated processor mode does not affect the traffic management and virtualization features of WebSphere Virtual Enterprise.

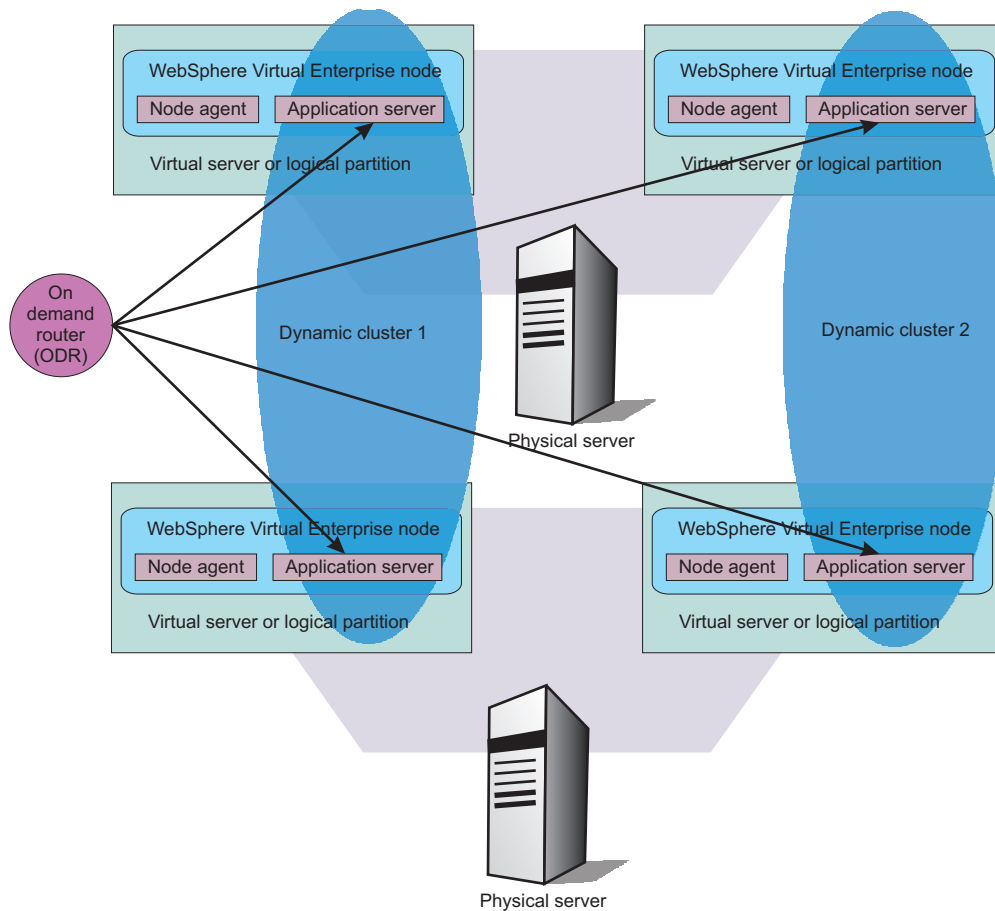


Figure 5. Coexistence of application infrastructure and server virtualization.

Chapter 2. Supported server virtualization environments

Before you deploy WebSphere Virtual Enterprise on virtualized servers, you must understand the limitations for the server platform that you are using.

Remember: This information changes frequently. Download the latest copy of this document on the WebSphere Virtual Enterprise Wiki or view the information in the WebSphere Extended Deployment information center.

Table 1. Server virtualization environments

Virtualization platform	Restrictions	Supported processor sharing mode
ESX Version 3.5 6.1.1+ VMware vSphere 4	<ul style="list-style-type: none"> See VMware Infrastructure 3 platforms and WebSphere Virtual Enterprise in the information center for a list of limitations. 	Shared and dedicated modes are supported for ESX or VMware vSphere 4, but you must configure WebSphere Virtual Enterprise to communicate with the ESX or VMware vSphere 4 hypervisor or the vCenter that is in control of the hypervisors where WebSphere Virtual Enterprise is running.
AIX® 5.3 and AIX 6.1 on POWER5™ and POWER6™ (Micro-Partitioning™)	<p>6.1.1+ See “Configuring AIX 5.3 and AIX 6.1 on POWER5 and POWER6 (Micro-Partitioning)” on page 11 for required configuration steps.</p> <p>For uncapped shared processor partitions, equal partition weights are recommended. WebSphere Virtual Enterprise does not have a partition to physical machine mapping and thus does not use the partition weight when making workload distribution and server placement decisions.</p>	<p>6.1.1+ All WebSphere Virtual Enterprise features are supported on AIX Micro-partitioning in capped or uncapped mode.</p> <p>Shared capped, shared uncapped, and dedicated modes are supported for both POWER5 and POWER6. Support for dedicated donating and multiple shared pools modes are available on POWER6 only.</p>
Linux® on POWER® (Micro-Partitioning)	None for dedicated.	Dedicated.

Table 1. Server virtualization environments (continued)

Virtualization platform	Restrictions	Supported processor sharing mode
Linux on z/VM®	<ul style="list-style-type: none"> • The guest operating system must be Red Hat Enterprise Linux (RHEL) 5.0, 5.1, 5.2 or SUSE Linux Enterprise Server (SLES) 10. • Mapping a Linux on zSeries® operating system image to a physical logical partition (LPAR): WebSphere Virtual Enterprise can balance workload across multiple z/VM virtual machines that are running Linux as the guest operating system. However, it does not have knowledge of the LPAR that is hosting the virtual machines on the z/VM image and is therefore unable to make workload balancing decisions based on the workload at the LPAR level. • Service policy goals: Workload that is not under the control of WebSphere Virtual Enterprise might be running on other Virtual Machines on the same z/VM image. This workload might affect service policy goals. 	Dedicated and shared.
Solaris Operating Environment 10 on Sun (Zones)	None for dedicated.	Dedicated.
HP-UX 11i v3 on HP using Virtual Partitions (vPars) or Integrity VMs	Currently not supported.	None.

Table 1. Server virtualization environments (continued)

Virtualization platform	Restrictions	Supported processor sharing mode
Linux Xen	<p>6.1.1+</p> <ul style="list-style-type: none"> • You must be using SLES 11 or later. • You must use the following guidelines for setting up the clock correctly: SUSE Linux Virtual Machines: Section 7.4.3 Virtual Machine Clock Settings. • Use para-virtualized mode over full-virtualized mode. Full-virtualized mode can lead to network bottlenecks that can adversely affect performance of the system. This performance problem is particularly possible for large topologies. 	
Microsoft® Hyper-V	Currently not supported.	None.

Chapter 3. WebSphere Virtual Enterprise in a micro-partitioned environment

You can use WebSphere Virtual Enterprise in a micro-partitioned environment. WebSphere Virtual Enterprise can understand the utilization of shared processor partitions.

Terminology for server virtualization on POWER AIX systems

You must understand the following terms when you are using WebSphere Virtual Enterprise on a POWER AIX system:

Logical partitioning

The ability to divide the resources of a system to create multiple separate servers. Each server runs its own operating system.

Micro-partitioning

The ability to share a pool of physical processors across multiple logical partitions. Physical processors can be allocated to a partition in increments of 0.01 of a processor with a minimum of 0.10 of a processor.

Shared processor partition

A partition that is configured to use a shared processor pool. A shared processor partition is a type of micro-partition.

Entitled capacity

The percentage of processor usage that is granted to a partition; specified in terms of 0.01 of a processor with a minimum of 0.10 of a processor.

Capped partition

A partition that cannot be granted more processing units beyond the configured entitlement for the partition.

Uncapped partition

A partition that can exceed its configured entitlement when needed, if resources permit.

Entitled capacity

A shared processor partition has a metric called **Entitled Capacity Percentage**. This metric represents the percentage the partition is using of its entitlement at a given point in time. The metric is visible in popular AIX system monitoring tools like the `lparstat` command, the `nmon` command, and the `topas` command.

An uncapped, shared processor partition can be assigned more processing capacity beyond its entitlement. The amount of processing capacity above its entitlement depends on the availability of processors in the shared pool and the maximum amount that the virtual processor configuration allows.

WebSphere Virtual Enterprise and shared processor partitions

WebSphere Virtual Enterprise Version 6.1.1 and later can understand the utilization of shared processor partitions and the dynamic capacity of the shared processor pool on the physical hardware that is required to operate in a shared processor partition environment.

WebSphere Virtual Enterprise operates in an uncapped, shared micro-partition environment by understanding two important metrics for each partition:

- Entitled capacity percentage
- Maximum entitled capacity percentage

WebSphere Virtual Enterprise calculates and uses these metrics so that the product can understand the utilization and dynamic capacity of a shared micro-partition.

WebSphere Virtual Enterprise and the entitled capacity percentage

The entitled capacity percentage represents the percentage the partition is using of its entitlement at a given point in time. This value represents the amount of processing power that is being used by that partition. WebSphere Virtual Enterprise publishes the value from each node agent or middleware agent process every 15 seconds for use by WebSphere Virtual Enterprise autonomic controllers such as the work profiler and dynamic workload manager (DWLM).

WebSphere Virtual Enterprise and the maximum entitled capacity percentage

When running on an uncapped, shared micro-partition, the maximum processing power of that partition is not fixed. In other words, the POWER Hypervisor™ can assign more processing power from the shared processing pool to a partition so that it is using more than its configured entitled capacity. For WebSphere Virtual Enterprise to understand the dynamic processor capacity, it calculates and publishes a metric called **maximum entitled capacity percentage**. This value represents the maximum percentage of its entitlement that a partition can use.

WebSphere Virtual Enterprise publishes the value from each node agent or middleware agent process every 15 seconds for use by WebSphere Virtual Enterprise autonomic controllers, such as the autonomic request flow manager (ARFM) and the application placement controller (APC). The maximum entitled capacity percentage value is calculated from the idle cycles of the shared processor pool to which the partition belongs. As the utilization of the shared processor pool increases and decreases, the maximum entitled capacity percentage value also changes.

WebSphere Virtual Enterprise metrics for shared processor partitions

WebSphere Virtual Enterprise can recognize additional statistics that are specific to uncapped, shared processor partitions. . These statistics can help you understand the utilization of a shared processor partition and the dynamic capacity of the shared processor pool. These statistics are helpful in monitoring WebSphere Virtual Enterprise when you are operating in a micro-partitioned environment.

Entitled capacity percentage

Represents the percentage of its entitlement that the partition is using.

Maximum entitled capacity percentage

Represents the maximum percentage of its entitlement that a partition can use.

Shared partition CPU utilization

Represents the processor utilization of an uncapped shared processor partition. The value of this statistic is calculated by the following formula:
(Entitled Capacity Percentage / Maximum Entitled Capacity percentage) x 100

Physical processors consumed

Represents the number of physical processors that are being used by the shared server partition.

Available pool processors

Represents the number of processors that are available in the shared processor pool to which this server partition belongs.

These statistics are saved in the NodeStatsHistoricCache log file. For more information about this log file, see NodeStatsHistoricCache topic in the WebSphere Virtual Enterprise Version 6.1.1 Information Center.

Configuring AIX 5.3 and AIX 6.1 on POWER5 and POWER6 (Micro-Partitioning)

To configure AIX Micro-Partitioning in shared and uncapped mode to work with WebSphere Virtual Enterprise, you must enable performance collection in the Hardware management console (HMC). You must enable this setting so that WebSphere Virtual Enterprise can recognize more than the entitled capacity for a logical partition (LPAR).

Before you begin

- You must have access to the HMC for your AIX servers.

About this task

If you do not configure this setting in your HMC, but you are using AIX micro-partitioning in shared and uncapped mode, the following message displays in your system out log of the node agent process:

```
ASPS0023W: The logical partition on which this node resides has a mode of shared, a type of uncapped, but performance collection is not enabled.
```

You must enable this setting so that WebSphere Virtual Enterprise can recognize more than the entitled capacity for a logical partition (LPAR).

1. Enable the performance hardware collection setting for each LPAR. In the HMC, click **Systems Management** → **Servers** → *physical_server_name* → *LPAR_name*. Select **Allow Performance Hardware collection**.
2. Click **OK**.
3. Restart the LPAR.

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, is the user's responsibility.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
500 Columbus Avenue
Thornwood, New York 10594 USA

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Mail Station P300
522 South Road
Poughkeepsie, NY 12601-5400
USA
Attention: Information Requests

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.