



VMware vSphere Online Library
ESX 4.0 Update 1 and Later
vCenter Server 4.0 Update 1 and Later

vSphere Resource Management Guide : Using NUMA Systems with ESX/ESXi : Specifying NUMA Controls

Specifying NUMA Controls

If you have applications that use a lot of memory or have a small number of virtual machines, you might want to optimize performance by specifying virtual machine CPU and memory placement explicitly.

This is useful if a virtual machine runs a memory-intensive workload, such as an in-memory database or a scientific computing application with a large data set. You might also want to optimize NUMA placements manually if the system workload is known to be simple and unchanging. For example, an eight-processor system running eight virtual machines with similar workloads is easy to optimize explicitly.

Note

In most situations, an ESX/ESXi host's automatic NUMA optimizations result in good performance.

ESX/ESXi provides two sets of controls for NUMA placement, so that administrators can control memory and processor placement of a virtual machine.

The vSphere Client allows you to specify two options.

CPU Affinity A virtual machine should use only the processors on a given node.

Memory Affinity The server should allocate memory only on the specified node.

If you set both options before a virtual machine starts, the virtual machine runs only on the selected node and all of its memory is allocated locally.

An administrator can also manually move a virtual machine to another node after the virtual machine has started running. In this case, the page migration rate of the virtual machine must be set manually, so that memory from the virtual machine's previous node can be moved to its new node.

Manual NUMA placement might interfere with the ESX/ESXi resource management algorithms, which try to give each virtual machine a fair share of the system's processor resources. For example, if ten virtual machines with processor-intensive workloads are manually placed on one node, and only two virtual machines are manually placed on another node, it is impossible for the system to give all twelve virtual machines equal shares of the system's resources.

Note

You can view NUMA configuration information in the Memory panel of the `resxtop` (or `esxtop`) utility.

Subtopics

[Associate Virtual Machines with a Single NUMA Node Using CPU Affinity](#)

[Associate Memory Allocations with a NUMA Node Using Memory Affinity](#)