



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](#) : Resource Management in NUMA Architectures

Resource Management in NUMA Architectures

You can perform resource management with different types of NUMA architecture. The systems that offer a NUMA platform to support industry-standard operating systems include those based on either AMD CPUs or the IBM Enterprise X-Architecture.

IBM Enterprise X-Architecture

One architecture that supports NUMA is the IBM Enterprise X-Architecture.

The IBM Enterprise X-Architecture supports servers with up to four nodes (also called CECs or SMP Expansion Complexes in IBM terminology). Each node can contain up to four Intel Xeon MP processors for a total of 16 CPUs. The next generation IBM eServer x445 uses an enhanced version of the Enterprise X-Architecture, and scales to eight nodes with up to four Xeon MP processors for a total of 32 CPUs. The third-generation IBM eServer x460 provides similar scalability but also supports 64-bit Xeon MP processors. The high scalability of all these systems stems from the Enterprise X-Architecture's NUMA design that is shared with IBM high end POWER4-based pSeries servers.

AMD Opteron-Based Systems

AMD Opteron-based systems, such as the HP ProLiant DL585 Server, also provide NUMA support.

The BIOS setting for node interleaving determines whether the system behaves more like a NUMA system or more like a Uniform Memory Architecture (UMA) system. See the HP ProLiant DL585 Server technology brief. See also the *HP ROM-Based Setup Utility User Guide* at the HP Web site.

By default, node interleaving is disabled, so each processor has its own memory. The BIOS builds a System Resource Allocation Table (SRAT), so the ESX/ESXi host detects the system as NUMA and applies NUMA optimizations. If you enable node interleaving (also known as interleaved memory), the BIOS does not build an SRAT, so the ESX/ESXi host does not detect the system as NUMA.

Currently shipping Opteron processors have up to four cores per socket. When node memory is enabled, the memory on the Opteron processors is divided such that each socket has some local memory, but memory for other sockets is remote. The single-core Opteron systems have a single processor per NUMA node and the dual-core Opteron systems have two processors for each NUMA node.

SMP virtual machines (having two virtual processors) cannot reside within a NUMA node that has a single core, such as the single-core Opteron processors. This also means they cannot be managed by the ESX/ESXi NUMA scheduler. Virtual machines that are not managed by the NUMA scheduler still run correctly. However, those virtual machines don't benefit from the ESX/ESXi NUMA optimizations. Uniprocessor virtual machines (with a single virtual processor) can reside within a single NUMA node and are managed by the ESX/ESXi NUMA scheduler.

Note

For small Opteron systems, NUMA rebalancing is now disabled by default to ensure scheduling fairness. Use the `Numa.RebalanceCoresTotal` and `Numa.RebalanceCoresNode` options to change this behavior.
