


VM Sizing Considerations

- A wide VM with more memory or vCPUs allocated than a single physical node contained for resources, would be node interleaved by the hypervisors, spreading the memory allocation and CPU scheduling across the physical nodes and removing NUMA locality




Virtual NUMA (vNUMA) Configuration

- Prior to ESX 5, if the number of vCPUs exceeded the physical cores in a single NUMA node, the VM would not be NUMA optimized by the hypervisor
- vNUMA in ESX 5+ and Windows Server 2012 Hyper-V allows larger VMs to be partitioned into virtual NUMA nodes to allow NUMA optimization of the VM in the hypervisor
- ESX 5+ only enables vNUMA when a VM has more than 8vCPUs or when the VM is wider than the number of cores in a physical node
 - Can be enabled manually by adding the *numa.vcpu.min* configuration parameter to VM configuration:
 - VM Settings->Options tab->Advanced General->Configuration Parameters
- vNUMA configuration in Windows Server 2012 Hyper-V is accomplished on the NUMA page for the VM settings

VM Sizing Considerations

- When a VM is smaller than a physical NUMA node for vCPU count and memory allocation, it is NUMA optimized by the hypervisor and all resources allocations occur from the same physical node for locality



vNUMA Optimized Configuration

- By configuring vNUMA for the wide VM, the vCPUs and memory allocations can be localized to the physical node configuration providing NUMA locality for scheduling and memory management to the VM
- For SQL Server this also creates a separate LazyWriter thread for buffer pool management under SQLOS

