virtensys

SOLUTION BRIEF

Solution: I/O Density: Enhancing x86 Server Virtualization

As server virtualization has become pervasive in today's data centers, if asked, most server administrators will admit that there are certain applications that are perceived as "harder to virtualize" than others. In fact, some application servers have been intentionally left off the server migration list, either because their requirements are deemed unsuitable for virtualization or their requirements (typically I/O requirements) are considered too resource intensive. There are many organizations that seek to become 100% virtualized within their data center. They have chosen to "*hyper-consolidate*" their server environment with a server virtualization platform such as VMware vSphere©.

However, a typical issue that server administrators face is when they have chosen to hyper-consolidate their server environment in an attempt to increase virtual to physical (V2P) ratios. They have usually chosen to leverage larger sized rack mount servers (2U or 4U form factor) simply because their I/O requirements force them to utilize a larger number of PCIe slots to house many independent I/O adapters. For example, after adding countless dual and/or quad port 1 Gigabit Ethernet (1 GbE) ports, they simply exhaust the physical confines of their server, their rack, and their network switch – from a cabling and port density perspective.

Thankfully, this can be remedied by providing enough I/O (and by choosing the right type and size of network interface) to be paired with their critical server hosts. In doing so, they can successfully virtualize a variety of server workloads, applications and/or desktop operating system images – by using Virtual Desktop Infrastructure (VDI). This greater resource potential and simplified management allows them to scale beyond the typical consolidation ratios, whilst meeting or exceeding the performance requirements of those "hard to virtualize" application servers.

USAGE SCENARIO

Let's look at a typical IT scenario:

Tabitha, the VMware administrator from XYZ Corporation, has been struggling with her company's server virtualization project. The finance department has asked for 10 new virtual machines to be created on her already overburdened vSphere server host. She has loaded up her physical server with a variety of VMs ranging from test/development to mission critical application servers.

www.virtensys.com

© Copyright 2011, Virtensys

Her existing network is a 1 Gigabit Ethernet (1 GbE) environment. Although, she has been planning with her network team to upgrade the segment where the vSphere hosts reside to 10 GbE, but they have yet to do so for a variety of issues: limited budget, not enough time, and not enough resources. Until then, Tabitha has make do with what she has. She has literally loaded up each of her standard rack-mount servers with as many 1 GbE network interfaces as the server can support. The back of her server looks like a "spaghetti factory".

She has configured several virtual switches, each using dedicated network interfaces. She has many network provisions to support network communications for server management, vMotion, storage and general application network traffic. Just as with her network environment, she has tapped out every available storage resource, using local disks as production VMs as well as her Tier 1 SAN. She has had little or no time to appropriately size or determine the best choice of storage for her very diverse virtual data center.

After adding the 10 additional new VMs to the test environment for her Finance department, some test users report that the new VMs and some existing virtual servers are running slowly. In fact, application response times are not as they had been prior to the addition of the new VMs. Tabitha suspects that she doesn't have enough network and storage I/O to satisfy the new workloads running on her vSphere server host. In addition, she is seeing disk I/O latency for a variety of storage operations. She decides to look for a better way to do things…

Is there a better way?

Yes, in this case --Tabitha decides to deploy PCIe sharing, delivered through Virtensys' VIO-4000 series. By doing so, she can now leverage a variety of high bandwidth, low latency network and storage interfaces. She can deploy dedicated 10 Gigabit Ethernet (10 GbE) network interfaces to each of her physical vSphere hosts in the data center -- this can be done with minimal complexity, minimal downtime and maximum investment protection.

Specifically, through the use of Virtensys' PCIe sharing technology, each server that is connected to the IOV appliance sees dedicated virtualized network connectivity. For example, with the Virtensys VIO-4001, using a single server connection, each server sees a pair of 10 GbE and 4/8 GB FC HBA interfaces. Instead of having multiple single, dual or quad-port 1 GbE network connections, Tabitha can now leverage 10 GbE and create multiple virtual networks and vSwitches to support a variety of network traffic types. Each can be associated with a different set of virtual 10 GbE interfaces and managed appropriately to the required service level of her applications.

In fact, with Virtensys' VIO-4000 series, administrators can apply "bandwidth control" at the vNIC level to guarantee a minimum percentage of bandwidth, allowing interfaces to always have the required amount but be able to burst to the full capabilities of the interface when needed.

In addition, Tabitha's diverse storage needs led her to leverage the Virtensys VIO-4008 and/or VIO-4010 as well, which allows standard rack mount servers to access virtual Direct Attached Storage (vDAS). Virtensys not only virtualizes network connectivity but also supports virtual RAID controllers (for presenting local SAS/SATA or SSD based disks, appearing as dedicated RAID sets and LUNs for each server). Tabitha can now utilize vDAS for her local datastores (removing the need for local disks within physical server) and the local boot vSphere partition.

© Copyright 2011, Virtensys

Virtensys' I/O Virtualization solution

Through the use of Virtensys' PCIe sharing appliance, the server administrator can reduce the physical complexity of the hypervisor host and allow for standardization for a more consistent deployment. "Disaggregating" the I/O resources and removing dependency on local disk resources means fewer points of failure, management, and simplified configuration when scale is needed most.

Ultimately, PCIe sharing provides a new way to deploy I/O to critical server hosts and to improve the communication between server hosts by providing dedicated 10 GbE and 8 GB Fibre Channel network/SAN network connections. For example, these can be ustilized to provide dedicated network connections to support: vMotion, Storage vMotion, Fault Tolerance, general application and management traffic.

Find out more:

Learn more about how Virtensys' PCIe sharing appliance enables you to virtualize your server workloads, applications and/or virtual desktop infrastructure; allows you to scale beyond "traditional" consolidation ratios and meet / exceed performance ratios.

Please visit us at: www.virtensys.com/solutions or email us at info_request@virtensys.com

© Copyright 2011, Virtensys