

How Xsigo Enhances VMware Deployments

Benefits of Server Virtualization

Virtualization technologies like VMware ESX Server play a critical role in boosting server utilization. Non-virtualized servers in the data center typically average very low utilization rates—less than 15 percent, according to Gartner. By enabling multiple operating systems and applications to run on a server simultaneously, OS virtualization drastically improves utilization.

Virtual machines also improve business continuity by eliminating the need to take applications offline when performing server maintenance. By running virtual machines on a hardware abstraction layer, you can easily migrate applications from one server to another with little or no disruption in service.

But the benefits come at a cost: Today's I/O infrastructure was not designed to run multiple applications on a single server. Nor was it designed to allow applications to migrate from server to server. The benefits of virtualizing the server's processors create new challenges that can best be resolved by virtualizing the servers' connectivity as well.

I/O Issues Created by Server Virtualization

Virtualization software like VMware ESX facilitates shared processor resources. But in this model another resource is shared as well: the server I/O. Because the I/O infrastructure was designed for traditional server usage, this new use case creates these I/O challenges:

I/O Bottlenecks: The dynamic nature of virtual machines and the greater utilization of server resources puts a significantly larger load on the servers' I/O. Higher traffic loads and more unpredictable traffic patterns may cause applications to slow down.

One possible remedy is to simply add more I/O: more Ethernet or FC connections. But I/O capabilities tend to be limited in today's 1U high or blade servers. These devices usually include just one or two PCI-Express slots, effectively limiting the total number of I/O ports to four Ethernet interfaces and two Fibre Channel interfaces. While this is adequate in traditional environments, in a server running four to twenty operating systems this connectivity can cause I/O to become a performance bottleneck.

Unpredictable performance: When multiple applications share a common set of resources, the performance of any given application may be unpredictable. A non-critical app may consume bandwidth or processor resources that are needed for a more critical task. While VMware ESX offers traffic shaping policies, these policies should be offloaded to external hardware

The security risk of "big flat networks": Today's storage and networking security techniques were not designed for the demands of virtual machines. Security today is controlled in several ways. The most obvious is the physical connection: don't connect a server where it does not need to be connected.

"Virtual machine mobility is driving the need for server I/O virtualization."

JOHN HUMPHREYS
PROGRAM VICE PRESIDENT
IDC ENTERPRISE
VIRTUALIZATION RESEARCH

Another is to control network and storage access at the I/O card level. The unique identifiers of the I/O cards are used to regulate connectivity by, for example, configuring a storage network to allow disk access to only a specific I/O card. That card has a specific worldwide name that is unique so no other card can access that storage. Hence, security is ensured.

This security breaks down if multiple VMs share one set of I/O cards. In that case, all VMs can potentially access the same storage resource. Furthermore, if we require the VMs to be mobile (ie, to have the ability to migrate through a VMotion event), the configuration must be “opened up” to grant access across multiple servers as well.

In the extreme case, suppose the deployment requires the flexibility to deploy any VM on any server. To accomplish this, the system must be configured with “flat networks” and totally open storage. In this case, all servers would be granted access to all resources, a model that compromises the basic concepts of security management.

Computer security expert A. Bryan Sartin of Cybertrust refers to “big flat networks” as one of the three most common security mistakes that companies make. Yet “big flat networks” are exactly what some companies are using today to implement virtual machines in large deployments.

How Xsigo Helps

The Xsigo I/O Director delivers virtual connectivity that addresses the I/O issues introduced by server virtualization. Benefits of virtual I/O include:

- Scalable connectivity to eliminate performance bottlenecks.
- Granular QoS for predictable I/O performance on specific virtual machines.
- Dedicated I/O resources to critical VM's ensures security.
- Portable I/O to maintain security as VMs are moved among servers.

Virtual I/O Replaces Physical I/O

Xsigo employs virtual I/O to deliver connectivity that is ideally suited to virtual servers. With Xsigo, the server's physical I/O cards (the NICs and HBAs) are replaced by virtual NICs and virtual HBAs. These virtual resources appear to the applications exactly as physical cards, but are far more flexible to manage. Benefits of virtual I/O include:

- Can deploy vNICs and vHBAs on the fly.
- No need for ESX reboot when connectivity is added.
- Hypervisors view virtual resource like as they view physical cards.
- Critical VMs can be associated with specific I/O resources for security.
- Virtual I/O can be migrated between physical servers.
- QoS controls can deliver guaranteed bandwidth to specific VMs.

These benefits bring significant advantages to the management of virtual machines.

Xsigo Eliminates I/O Bottlenecks

Xsigo eliminates bottlenecks by delivering more bandwidth to the server and allowing that increased bandwidth to be better used by VMs.

More bandwidth: Xsigo delivers 10Gb to each server, through a single pipe that can be dynamically allocated across VMs and across both storage and networking traffic. If network traffic needs 10Gb at this moment, the full bandwidth is there. If storage needs it the next moment, the full 10Gb is there. Conventional connectivity offers less throughput and less flexibility.

Better usage by VMs: Virtual NICs and virtual HBAs can be assigned through ESX to individual virtual machines for complete control over resources. Better yet, at any given time any of the VMs can have access to the full 10Gb bandwidth of the external pipe, if needed.

Xsigo Enhances Security for Virtual Machines

Security is enhanced by associating virtual connectivity with specific VMs. This level of management granularity enables security management using the same best practices that apply to conventional servers.

Control via connectivity: The simplest security device is connectivity. If the VMs on a particular server do not need access to a network, simply do not associate connectivity. For example, the VMs on a server can be protected from the public network by simply not associating the server with a vNIC on that network. At the same time, a different server can be granted access to the public network without compromising security on the first server.

Control via identity: Associating a virtual I/O resource with a VM makes it simple to manage security. Each virtual resource has a unique identity (a WWN or MAC address) that enables all network and storage security measures to retain their full effectiveness. Furthermore, the I/O resource and its identity can be moved among servers along with a VM. This means the IT manager is never forced to implement “big flat networks” or open storage.

Xsigo Helps Deliver Predictable Performance

Quality-of-service profiles can be configured per virtual I/O resource, which makes it possible to control I/O flows per VM, if each VM has its own dedicated virtual I/O resources.

For each vNIC or vHBA, a committed information rate (CIR) guarantees a minimum bandwidth, and a peak information rate (PIR) limits the maximum amount of bandwidth the resource can consume. By properly tuning these settings for the applications running on the virtual machine, IT managers can ensure predictable network performance for critical applications.

Summary

The virtual I/O capabilities of the Xsigo I/O Director enable a virtualized datacenter running VMware ESX to scale better, ensure security through traffic isolation, and guarantee performance through QoS. Together, these benefits allow a greater percentage of applications to run in virtual machines, further reducing the need for physical servers that consume power and real estate.

Benefits of VM Live Migration

Virtualization technologies like VMware ESX Server play a critical role in boosting server utilization. Tools like VMotion, which allow virtual machines to be easily migrated from one server to another, help to further increase datacenter utilization. By treating servers like a flexible pool of resources and easily moving VMs from one server to another, the datacenter can quickly respond to load changes, rapidly address system failures, and more effectively meet the needs of critical applications.

VMotion is able to migrate live virtual machines (VM's) with no impact on the users by transmitting the active memory and execution state of a virtual machine from one instance of ESX to another. The VM's stored data or disk is accessed by the new server through shared storage using VMFS cluster file system. Once the new server has the same files, memory and execution state as the old server, the VM and its applications begin running seamlessly on the new server.

VMotion Network Design

The live migration of VM's taxes the network infrastructure in ways that other traffic does not. In order to guarantee that VMotion runs as smoothly as possible, these critical aspects of network design must be considered:

- 1) **A high-speed network is required.** Because VMotion needs to transmit large amounts of memory state as quickly as possible, a high-bandwidth network is needed.
- 2) **The network should be congestion free.** Congestion limits the amount of bandwidth available to VMotion transfers.
- 3) **The network should be isolated.** VMotion transmits the active memory and execution state information in the clear. Consequently, to ensure that private information remains private, all VMotion traffic should be on an isolated network.

If the VMotion network is not able to transfer the memory and execution state quickly enough, the VMotion event may take longer than acceptable or may fail altogether. If the memory and execution state of a VM changes faster than VMotion can transfer the data, the VM will never be able to be migrated. Fortunately, Xsigo's I/O Virtualization can easily address these issues and help ensure that VMotion runs smoothly.

Xsigo I/O Virtualization and VMotion

By using the Xsigo VP780 I/O Director as the infrastructure for deploying virtual machines, a VMotion network can be configured that is both extremely fast and isolated. Xsigo virtual NICs and a dedicated Ethernet I/O Module can create an internally-switched, high-speed, low-latency network that allows VMotion to transfer data over an isolated Ethernet network at rates up to 7.4 Gbps. Alternatively, using InfiniBand interfaces on Xsigo's 780Gbps InfiniBand fabric to migrate VMs enables extremely fast transfer rates and isolation from the LAN.

