THE TECH INSIDER'S GUIDE TO I/O VIRTUALIZATION

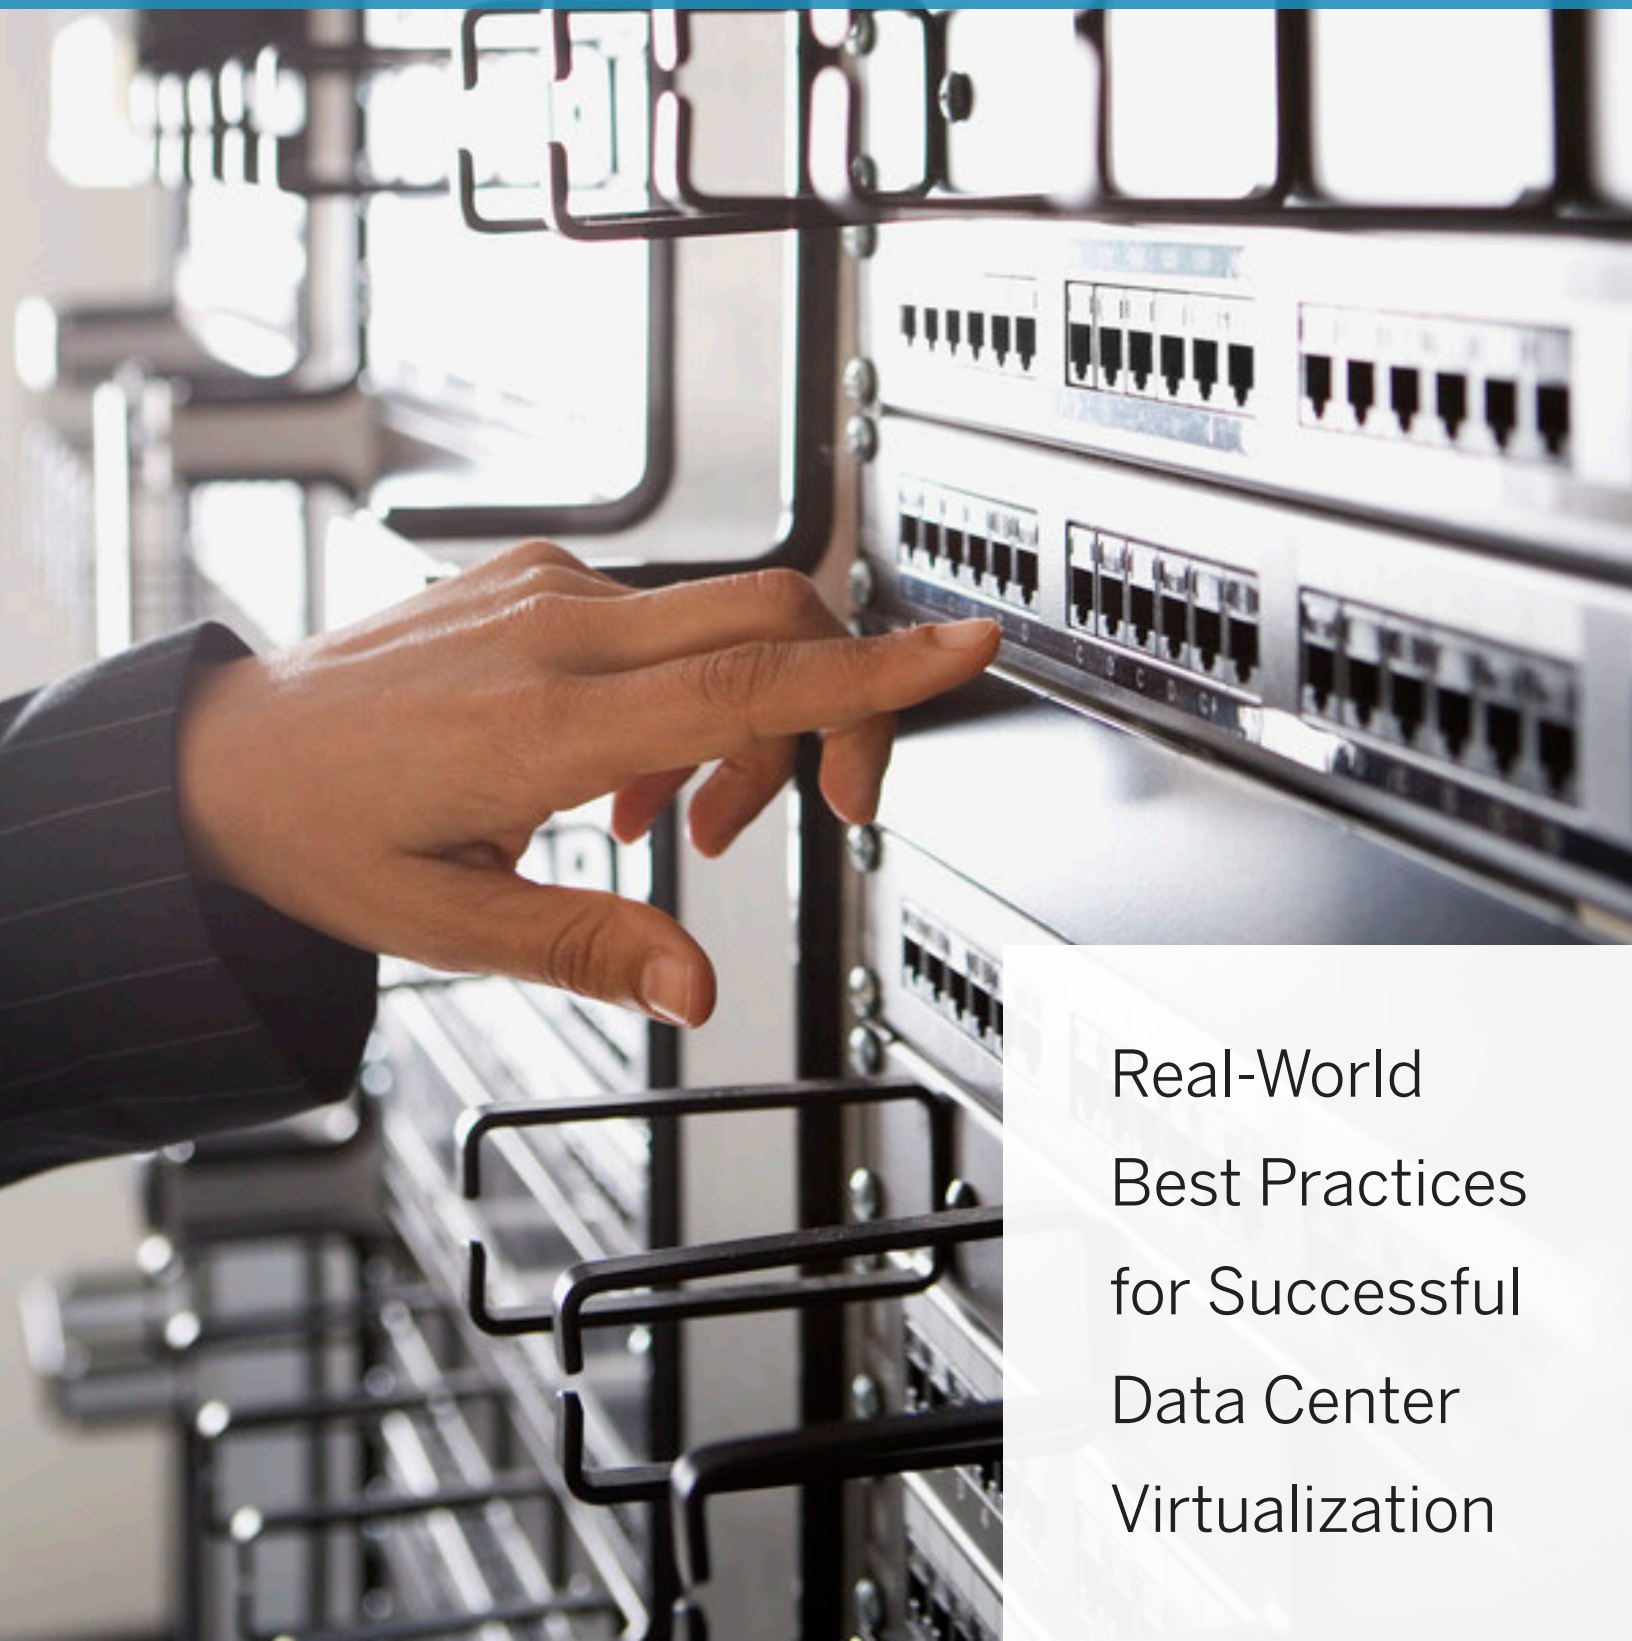

Real-World Best Practices for Successful Data Center Virtualization

THE TECH INSIDER'S GUIDE TO I/O VIRTUALIZATION

The Virtual I/O Revolution

Virtual I/O Overview

I/O Virtualization at a Glance

Xsigo I/O Virtualization Solutions

The Virtual I/O Revolution

In today's leading-edge data centers, virtualization technologies boost the utilization of storage, server, and networking resources, allowing these assets to be deployed with much greater speed and flexibility than ever before.

On the other hand, server connectivity remains much as it was ten years ago. IT managers must still contend with a maze of cables, cards, switches and routers—and the management complexity that accompanies them. This inflexible I/O infrastructure drives up cost and slows resource deployment, making it impossible to accommodate changing business requirements in real time.

Virtual I/O changes the game. By replacing fixed I/O cards with virtual I/O resources, IT managers can significantly enhance data center agility. Tasks that formerly took weeks can now be done in minutes. Because connectivity is consolidated, the infrastructure becomes dramatically simpler: Hundreds of cables are replaced by dozens, most I/O adapter cards are eliminated, and overall connectivity costs drop by up to 50%.

This guide reviews the factors driving virtual I/O adoption, how virtual machine technologies benefit from virtual I/O, and the specific solutions that are available now.

The Problem with Today's I/O

Server I/O is a major obstacle to effective server virtualization. Virtualized servers demand more network bandwidth, and they need connections to more networks and storage to accommodate the multiple applications they host. VMware best practices recommend dedicated connectivity for critical virtual machines and management networks. According to a recent survey of virtualization users, a typical virtualized server will include seven or more I/O connections, with some users reporting as many as 16 connections per server. Furthermore, a virtualized server is far more likely to need I/O reconfiguration as requirements change.

Virtualization introduces performance questions as well. When all I/O resources are shared, how do you ensure the performance and security of a specific application?

For these reasons, IT managers are looking for new answers to the I/O question.

"While blade server hardware and virtualization software enhance data center management, today's server I/O still limits agility. I/O virtualization promises to deliver the next significant advancement. Just as server virtualization decouples applications from specific servers, virtual I/O will accelerate application management by removing the constraints of the fixed I/O infrastructure."

DR. BERNARD S. MEYERSON

IBM FELLOW
VP OF STRATEGIC ALLIANCES
CTO

IBM

Why Virtualized Servers are Different

Traditional servers do not necessarily encounter the same I/O issues as virtualized servers. Traditional servers are often deployed in the three-tier data center model (shown below) where servers are configured with the I/O necessary to meet that device's specific application requirements. Database servers, for example, will connect only to the SAN and to the application tier. Servers in the application tier connect only to the database and Web tiers. These physically distinct networks provide security and availability while isolating devices from intrusion threats, denial of service attacks, or application failures.

Server virtualization completely changes the deployment model. The objective of virtualization is to increase server utilization by creating a dynamic pool of compute resources that can be deployed as needed. There are several requirements to achieve this objective:

- **Flexible application deployment:** Most applications should be able to run on any server.
- **Multiple applications per server:** If a server is underutilized, applications can be added to capitalize on available compute resources.
- **Application mobility:** Applications should be portable among all servers for high availability or load balancing purposes.

These requirements have important implications for server I/O:

- **Increased demand for connectivity:** For flexible deployment, servers need connectivity to all networks. A server may connect to the SAN and to the DMZ, and to every network in between. To ensure isolation, many users require separate physical connections to these networks, which results in numerous physical connections to each server.

Traditional three-tier data center architecture.

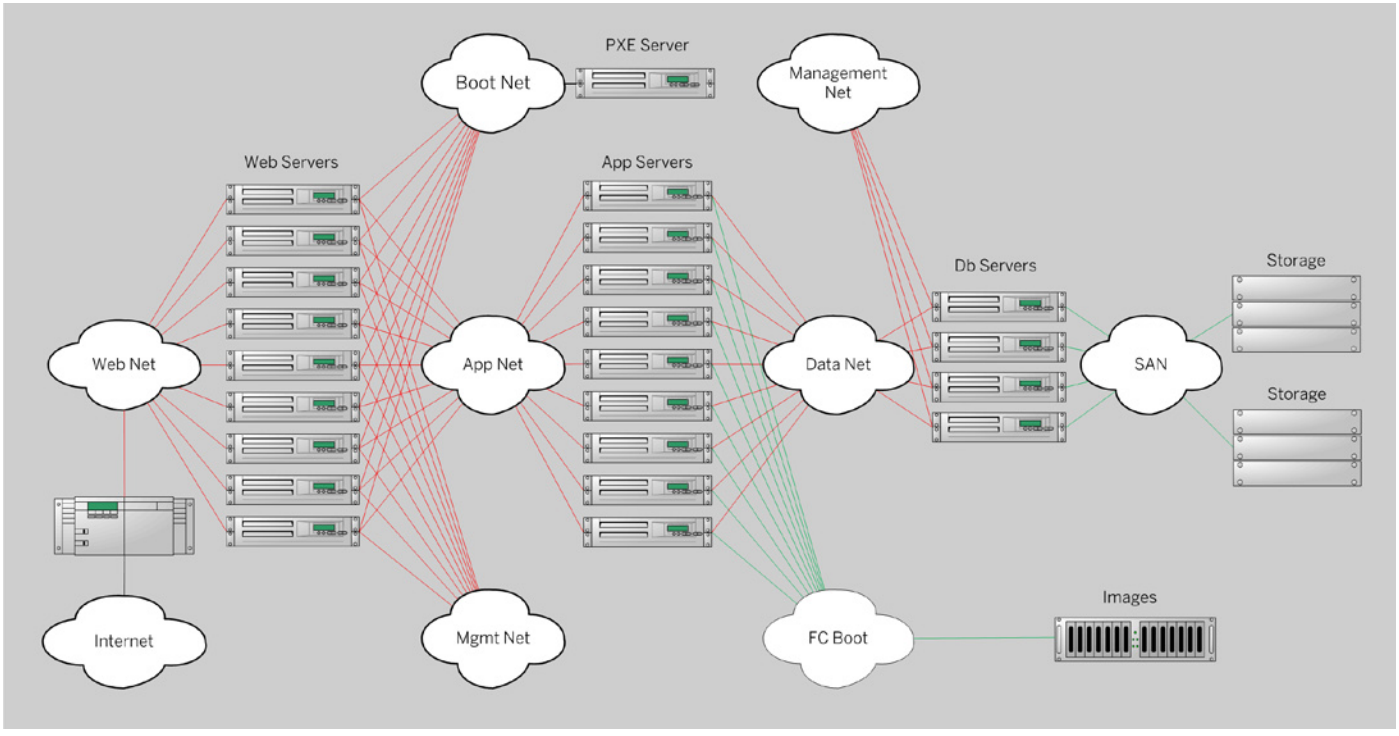

- **Management networks:** Virtualization management places additional demands on connectivity. VMware best practices, for example, dictate dedicated connectivity for the VMotion network.
- **Bandwidth requirements:** Virtualization increases the bandwidth demands of server I/O. Traditional servers often operate at only 10% processor utilization; virtualized devices may grow beyond 50%. A similar increase will occur in I/O utilization, revealing new bottlenecks in the I/O path.

No Attractive Options with Traditional I/O

Traditional I/O, designed for one-application-per-server deployments, can be repurposed for virtualized servers. But the options have limitations.

- **Use the server's pre-existing I/O:** One option is to employ conventional I/O—a few NICs and perhaps a few HBAs—and share it among virtual machines. For several reasons, this is not likely to work. Applications may cause congestion during periods of peak performance, such as when VM backups occur. Performance issues are difficult to resolve since diagnostic tools for VMware are still relatively immature and not widely available. Even if an administrator identifies a bottleneck's source, corrective action may require purchasing more network cards or rebalancing application workloads on the VMs.

  Another issue with pre-existing I/O is the sheer number of connections needed with virtualization. Beyond the numerous connections to data networks, virtualized servers also require dedicated interconnects for management and virtual machine migration. Servers are also likely to require connectivity to external storage. If your shop uses SAN, this means FC cards in every server.

Efficient virtualized data center architecture with virtual I/O.

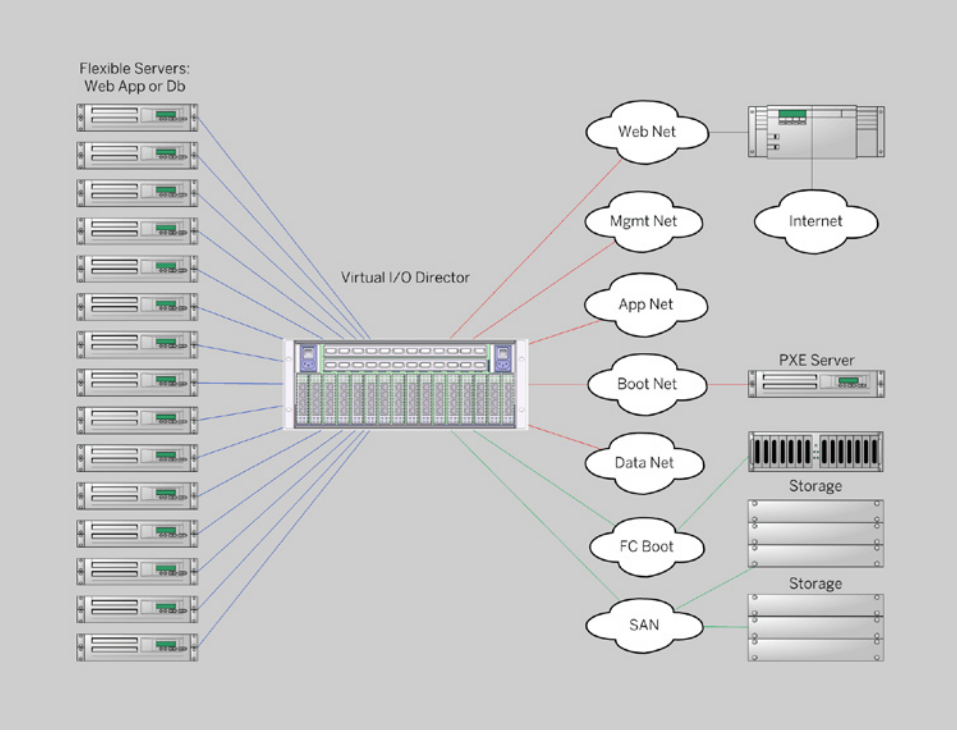

- **Scale up the I/O for each server:** Most IT managers increase connectivity to accommodate virtualization. While this is a viable option, it is not always an attractive one. Virtualization users find they need anywhere from six to 16 I/O connections per server, which adds significant cost and cabling. More importantly, it also drives the use of larger servers to accommodate the needed I/O cards, which add cost, space and power requirements. The end result is that the cost of I/O can frequently exceed the cost of the server itself.

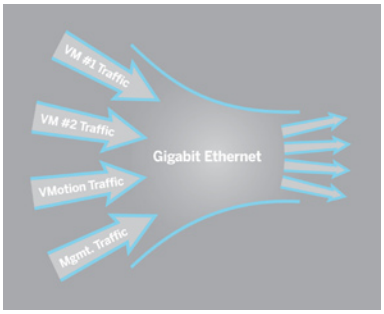

Traffic congestion on a gigabit ethernet link.

Blade Servers Solve One Problem, Create Another

Server blades solve the cabling issue with an internal backplane that reduces connectivity issues. Limits on the number of connections, however, can be problematic for virtualization. Accommodating high levels of connectivity may prove to be either costly (sometimes requiring an expansion blade that consumes an extra slot), or impossible depending on the requirements.

Virtual I/O Overview

Virtual I/O presents an option that is scalable, flexible, and easily managed. It removes the limitations of traditional I/O and provides optimal connectivity when running multiple VMs on a single physical server:

- **I/O Consolidation:** Virtual I/O consolidates all storage and network connections to a single cable per server (or two for redundancy). This reduces cabling and allows additional "connections" to be added whenever required. The single link must have these characteristics to fully consolidate connectivity:
- **Fibre Channel-ready:** Virtual I/O carries Fibre Channel traffic for connection to FC storage.
- **High Performance:** Virtual I/O has sufficient performance to meet all of the server's aggregate I/O requirements. The single link will never be a bottleneck.
- **Quality of Service:** Virtual I/O's QoS controls manage bandwidth allocation, which helps manage application performance in a shared resource environment.
- **Transparent I/O Management:** Virtual I/O appears as conventional connectivity resources (NICs and HBAs) to the application, OS, and hypervisor. It accomplishes this with virtual NICs (vNICs) and virtual HBAs (vHBAs). Much like "virtual machines" appear to applications as distinct servers, vNICs and vHBAs appear exactly as their physical counterparts would, which allows them to be easily and broadly deployed.
- **I/O Isolation:** From a security standpoint, vNICs and vHBAs provide the same isolation as if they were physically separate adapter cards connected to separate cables.

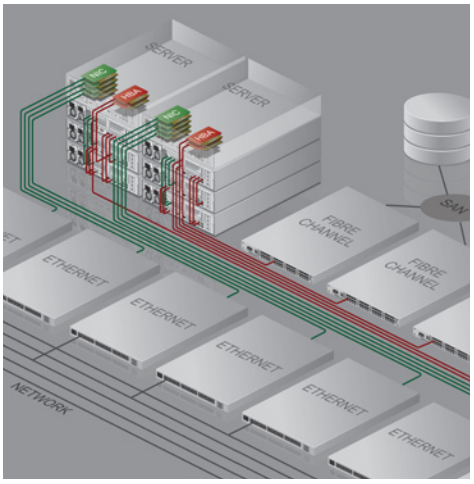

Data center without virtual I/O.

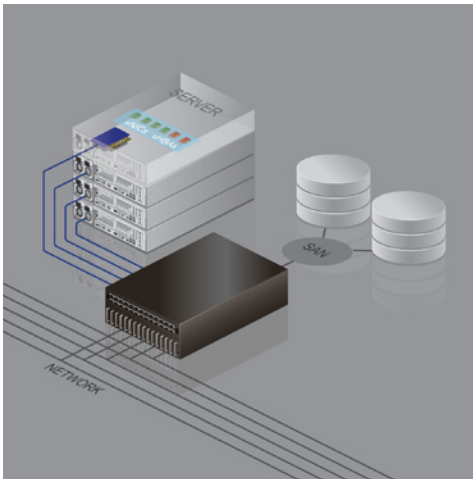

Data center with virtual I/O.

Elements of Virtual I/O

Virtual I/O maintains the appearance of traditional I/O (from the standpoint of servers and networks) while eliminating or reducing the need for components like adapter cards, cables, and edge switches. With virtual I/O, system elements are built specifically for the I/O virtualization task.

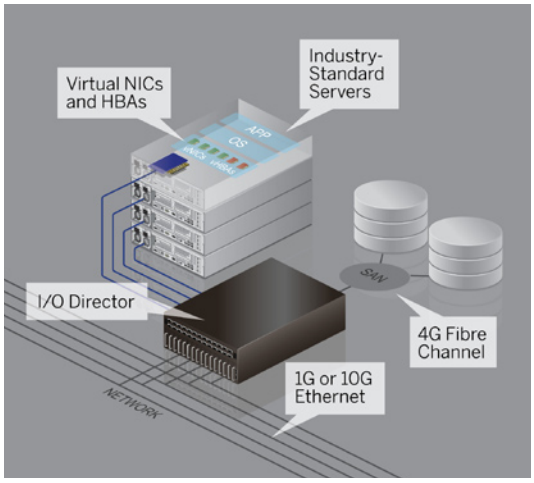

Virtual I/O in the data center.

- **Virtual I/O director:** Application servers connect to this top-of-rack device via individual links. The I/O director then connects to core Ethernet and SAN devices via conventional Ethernet and FC links.
- **High-speed interconnect:** Each server connects to the I/O director via a dedicated link (for redundancy, two I/O directors may connect to each server by separate links). Because this link must accommodate both FC and Ethernet traffic, InfiniBand is today the only transport that meets the requirements.
- **Host Channel Adapter (HCA):** A single card in each server (or two for redundancy) provides connectivity to the I/O director(s).
- **Virtual NIC (vNIC):** Appears exactly as a physical network adapter to the application, operating system, and hypervisor. Capabilities of the vNIC include:

  Deployable on demand: vNICs can be created on the fly (no need for re-boot).

  Scalable: Multiple vNICs can be created on a single server.

  Mobile: vNICs can be moved among physical servers under the control of either the administrator or a management application.

  Persistent identity: MAC addresses remain persistent, even through re-boots or migration. This allows server I/O re-configuration without affecting network settings.

- **Virtual HBAs (vHBA):** Like the vNIC, the vHBA behaves exactly as a physical Fibre Channel host adapter. It can be deployed on demand, is scalable, mobile, and has a persistent WWN.
- **Quality of Service (QoS):** User-defined QoS attributes define guaranteed performance parameters for individual vNICs and vHBAs.

What I/O Virtualization is Not

Virtual I/O is designed specifically to meet the needs of server I/O. It neither conflicts nor overlaps with server or storage virtualization. It is in fact, fully complementary with those and other management solutions.

- **Virtual I/O is not a hypervisor:** Virtual I/O provides network and storage connectivity for applications, operating systems, and hypervisors. It does not virtualize servers.
- **Virtual I/O is not storage virtualization:** Virtual I/O presents storage to servers exactly as a conventional HBA does. There is no re-mapping of LUNs.

Why InfiniBand is an Ideal Interconnect for Virtual I/O

InfiniBand (IB) was developed specifically as an interconnect technology for converged I/O fabrics. The requirements for these fabrics are:

- **Reliable transmission with flow control:** IB provides reliable link and transport layers with flow control, all implemented in hardware.
- **High performance:** IB Latency is very low: 1usec at the endpoints, 150ns at switch hops. Throughput is high: 10Gbs and 20Gbs. Data rates are non-available.
- **Efficient data transfer:** IB provides RDMA which enables zero-copy transfers, thus eliminating the need to waste CPU cycles and memory bandwidth on copying data.
- **Dedicated resources for different traffic flows:** IB provides virtual lanes with their own resources (such as switch buffers) for different traffic flows (such as networking and storage).
- **Scalability:** IB scalability to thousands of nodes has been proven in large clusters—a quarter of the TOP500 clusters use it as their interconnect.

InfiniBand delivers in all these critical areas for I/O fabrics, and it is now a mature technology which is offered as an option by major server vendors. Operating systems with mature IB support include Linux and Windows. The technology has been proven in many large clusters, including ones running mission-critical applications such as high-frequency trading. IB provides an ideal technology for consolidating data and storage traffic by offering the best transport currently available for I/O virtualization.

How I/O Virtualization Enhances Server Virtualization

Virtual I/O to delivers connectivity that is ideally suited to virtual servers. The two are very much analogous: While server virtualization dynamically allocates high-speed processing resources among multiple applications, I/O virtualization dynamically allocates high-speed I/O resources among multiple demands. From the application perspective, both are managed as traditional resources.

Benefits of virtual I/O include:

- **Agility:** You can deploy vNICs and vHBAs to virtualized servers, without server reboots.
- **Consolidation:** With one cable, you can obtain the multiple connections needed for VMs and management.
- **Continuity:** Hypervisors view virtual resource exactly as they view physical cards.
- **Resource isolation:** VMs can be associated with specific I/O resources for security. The simplest security device is connectivity. If a particular server does not require access to a particular network, it is a simple, software controlled process to remove that connectivity.
- **Mobility:** Virtual I/O can be migrated between physical servers to assist with VM migration.
- **Predictability:** Quality of Service delivers guaranteed bandwidth to specific VMs.

Six use cases show how these attributes enhance virtual server deployments.

1. Predictable Application Performance

It can be challenging to guarantee application performance and isolation when relying on shared resources. In production deployments, it's important to eliminate bottlenecks and guarantee throughput for critical applications.

With traditional I/O, IT managers resolve these issues by adding resources. Multiple connections to each server provide isolation and bandwidth, but they also drive up costs to the point that server I/O may be more expensive than the server itself.

Virtual I/O provides a less costly, more easily managed alternative that delivers predictable application performance in three ways:

- **Dynamically allocated bandwidth:** 10-20Gb bandwidth to each server is dynamically shared among virtual machines. Conventional I/O may get only 1Gb bandwidth to an app. With virtual I/O, all bandwidth is available when needed.
- **Resource isolation:** Connectivity can be assigned to specific VMs for I/O isolation.
- **Quality of Service (QoS):** Ensures that critical applications receive the bandwidth required.

QoS helps ensure predictable application performance by delivering guaranteed bandwidth, even when resource contention occurs. Throughput is hardware enforced, and controlled by these user-defined settings:

- Committed information rate (CIR) guarantees a minimum bandwidth.
- Peak information rate (PIR) limits the maximum amount of bandwidth the resource can consume.

QoS settings can be managed in two ways:

- **By virtual machine:** QoS can be set by virtual machine, regardless of which vNIC or vHBA the VM is associated with. If the VM is moved to another server, the QoS setting remains with it. This powerful set-it-and-forget-it feature allows QoS management through VMotion events.
- **By virtual I/O resource:** QoS set per virtual NIC or HBA. When the resource is migrated to another server, the QoS settings remain intact.

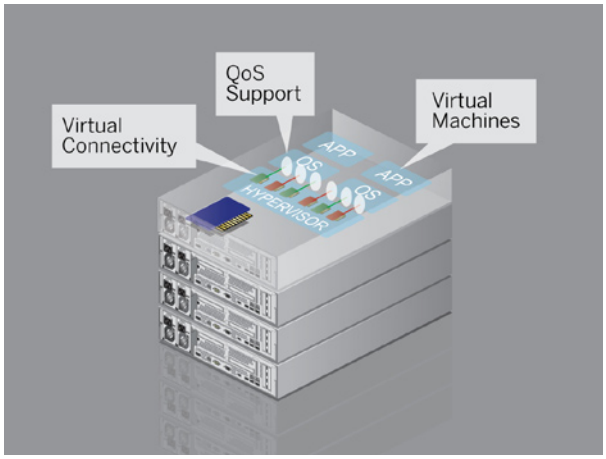

Virtual I/O enables QoS for each virtual machine.

In addition to isolating traffic, virtual I/O allows connectivity to be added based on need. This capability allows a virtualized datacenter to be both flexible and secure. For example, if only a few of the virtual machines require access to the internet through the firewall DMZ, only the ESX servers hosting those virtual machines need connectivity to the firewall DMZ. By removing connectivity to the firewall DMZ, a potential vulnerability is eliminated from the servers. When demand for more DMZ connectivity occurs, the virtual I/O solution can establish that connectivity for additional ESX servers.

Isolating traffic and connectivity improves deterministic performance of virtual machines and eliminates potential security risks that are prevalent in flat network architectures.

2. Lower Cost and Power

With conventional I/O, the connectivity requirements for virtualization often drive deployment of costly 4U servers. 1U servers may cost a fraction as much, but are typically limited to two Fibre Channel and two Ethernet interfaces.

By eliminating connectivity constraints, virtual I/O enables 1U servers to be more widely deployed. With just two adapter cards, as many as 32 fully redundant vNICs and 12 redundant vHBAs may be deployed to a server. Fewer I/O cards means less power as well, saving as much as 40 watts per server.

Edge switches and cables are also reduced, for a total capital savings of 30% to 50% on server I/O costs alone. Power savings associated with the reduced number of PCI cards and networking switches range from 4 to 6 kilowatts for a 120 server installation.

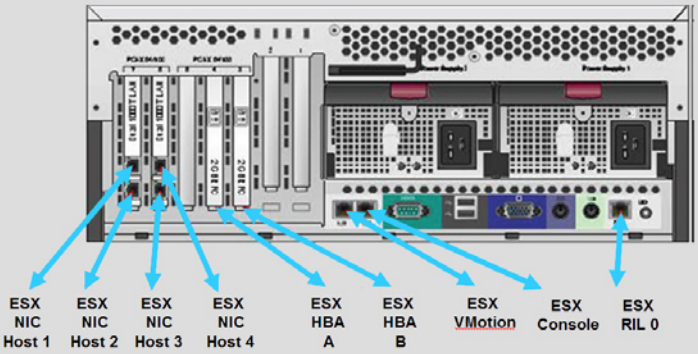

With conventional I/O, more connectivity often requires 4U servers.

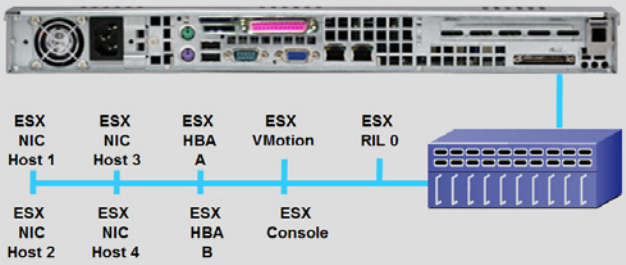

With virtual I/O, 1U servers can be used.

3. Blade Server Enablement

Blade servers, like 1U rack mount servers, typically have very limited I/O expansion capabilities. In many cases, a fully populated server blade can only have four Ethernet NICs and two fibre channel HBAs. This limited I/O capability significantly reduces the capability of blade environments to run multiple virtual machines per host. Sharing I/O resources across many virtual machines will result in conflicts and bottlenecks that could be avoided if more I/O resources were available.

Virtual I/O resolves the blade server I/O bottleneck. By replacing the traditional I/O resources with high-speed links and virtual resources, the virtual machines receive more bandwidth and better isolation. Most blade systems are available with InfiniBand cards that are proven compatible with virtual I/O.

4. I/O Resource Mobility

Virtual I/O eases the challenge of building fault-tolerant datacenter environments by enabling virtual I/O resources to be rapidly moved from one server to another. Because a server's I/O personality is typically defined by its connectivity, MAC addresses and WWNs, moving the virtual I/O resources effectively moves the server personality. This enables the physical resource running a hypervisor to be quickly changed from one server to another in the event of a system failure or system upgrade.

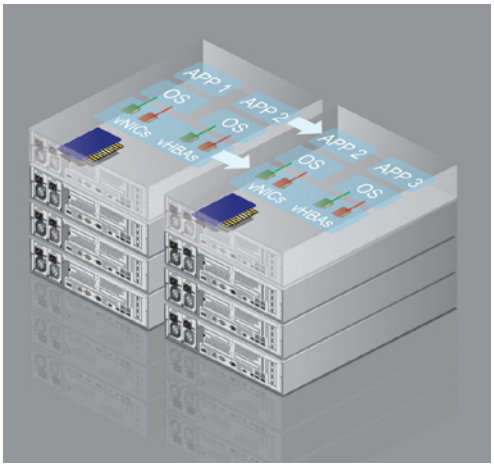

Server profiles can move from machine to machine.

The I/O mobility inherent in virtual I/O helps enable:

- **Disaster recovery:** The exact I/O resources at one site can quickly be created at another, thus accelerating switchover.
- **Server replacement:** When a server fails, its full connectivity profile (including WWNs and MAC addresses) can be instantly moved to another server.
- **Security:** To ensure security, a server may be configured only with the connectivity currently needed. When requirements change, the configuration can be changed without entering the data center and the IT manager is never forced to implement "big flat networks" or open storage.

5. Secure VMotion

In most VMware deployments, VMotion information travels over the same Ethernet switching infrastructure as production traffic. This creates a security exposure that can be eliminated with virtual I/O. By using the high-speed, low-latency fabric as the virtual machine infrastructure, VMotion traffic remains on a network that is both extremely fast and separate from production traffic. 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 10 Gbps.

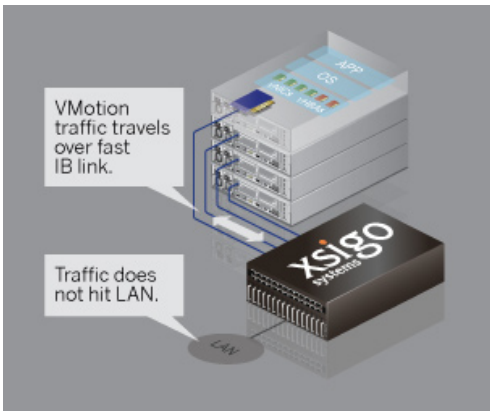

Virtual I/O ensures VMotion security.

- Improves security by isolating VMotion traffic from the rest of the datacenter
- Improves VMotion performance and reliability by creating a high-speed network for VMotion traffic
- Reduces the burden on the LAN by keeping all of the VMotion traffic within the Xsigo network
- Reduces the need for expensive high-speed networking ports and server NICs

6. Virtual Machine Migration Across Many Servers

Security requirements sometimes limit the use of VMotion. To migrate a VM from one server to another, VMotion requires that both servers simultaneously "see" the same networks and storage. All servers participating in the VMotion cluster would therefore need "open" access, a configuration that violates security requirements in some IT shops.

Virtual I/O enables an alternative to VMotion that allows virtual machines to migrate without the need for open access. By migrating I/O, a VM can be suspended on one server, then immediately resumed on another, without ever exposing storage to two servers simultaneously. This is accomplished in three steps:

- Suspend the virtual machine on Server A
- Migrate the applications storage and network connectivity (vHBAs and vNICs) to Server B.
- Resume the virtual machine on Server B.

This simple process works because the vHBA maintains its WWN through the migration. The target LUNs and associated zoning are therefore migrated with the vHBA which allows the server administrator to resume the virtual machine on the new server without any storage configuration changes.

I/O Virtualization at a Glance

When deployed with server virtualization, I/O virtualization provides several operational efficiencies that save cost and time.

Infrastructure Benefits

Attribute	Virtual I/O	Physical I/O
Cables per server	2 cables	6 to 16 cables
Redundancy on every I/O connection	Yes	No, for most installations
Minumum server size	1U	2U or 4U
Minumum server cost	$2K	$10K
I/O expense per server (including switch ports)	$2-3K	$4 - 6K
Quality of Service (QoS) management per interconnect	QoS per virtual NIC or HBA, enforced in hardware. I/O resources and bandwidth managed from I/O Director.	QoS is managed either: ▪ In the server (software consumes server resources). ▪ At the switch (requires switch that supports QoS). ▪ Neither provides central management of I/O and performance.
Quality of Service (QoS) management per virtual machine	QoS per virtual machine enforced in hardware. QoS attributes remain with VM through migration.	QoS enforced per connection only.
VMotion interconnect	Communicate directly from one vNIC to another via InfiniBand fabric. No LAN traffic.	Dedicated Ethernet port in each server. Connect via switch.

Management Efficiency

Task	Virtual I/O	Physical I/O
Add network connectivity (NICs)	1. Create a vNIC in each server. **Time: 5 minutes**	1. Add an Ethernet port to every server. 2. Add cabling. Add switch ports if needed. **Time: Days**
Add FC storage connectivity (HBAs)	1. Create a vHBA in each server. **Time: 5 minutes**	2. Add an FC HBA card to every server. 3. Add cabling. 4. Add FC switch ports if needed. **Time: Days**
Add a server	1. Install 1 or 2 host cards in server. 2. Rack server. 3. Connect 1 or 2 cables. 4. Deploy vNICs and vHBAs. **Time: 2-4 hours**	1. Install 2 to 6 host cards in server. 2. Rack server. 3. Add Ethernet and FC switches if needed. 4. Connect 6 to 16 cables 5. Individually configure NICs and HBAs. **Time: 8 hr+**
Failover VM instance from one server to another	1. Migrate connectivity. 2. Restart new server. VMs come up automatically. **Time: 10 minutes**	3. Add NICs to new server, if needed. 4. Configure NICs 5. Remap networks for new MAC address. 6. Add HBAs to new server, if needed. 7. Configure HBAs. 8. Remap networks to reflect new WWNs. 9. Restart server. **Time: 4 hours**

Xsigo I/O Virtualization Solutions

Xsigo Systems technology delivers on the promise of I/O virtualization. The Xsigo I/O Director™ is an enterprise-class, hardware-based solution that provides LAN and SAN connectivity for up to hundreds of servers. Under software control, any server may be configured with access to any LAN or SAN. Because the vNICs and vHBAs can be configured on-the-fly without a server reboot, the server's network connectivity is truly flexible and reconfigurable. Xsigo technology includes:

The Xsigo I/O Director chassis: Includes 24 InfiniBand ports (10-20Gb), 15 I/O module slots, and fully redundant power and cooling. Directly connect up to 24 servers, or connect hundreds of servers via expansion switches.

The Xsigo VP780 I/O Director

I/O Modules: Provide connectivity to LAN and SAN. Options include 1Gb and 10Gb Ethernet, and 4Gb FC.

XMS Management System: Provides central management of all I/O resources via CLI, GUI, or an XML-based API.

System Management

Xsigo Management System (XMS) is a powerful, easy-to-use, browser-based interface. Wizards guide the user through the process of creating server I/O profiles, individual virtual resources or templates. XMS also provides reporting views which enable an administrator to see all of the servers, virtual NICs or virtual HBAs that are currently configured. Traffic monitoring tools also enable administrators to understand how resources are being used.

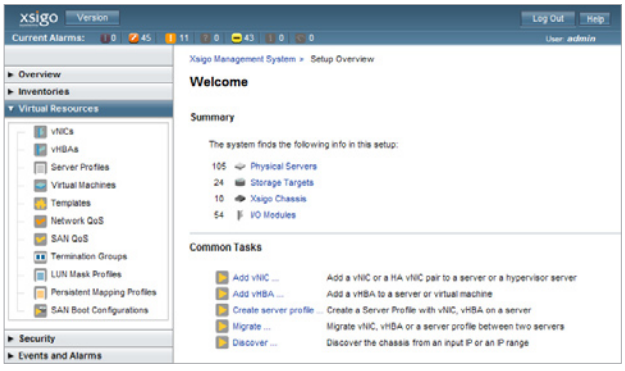

All I/O administration can be handled from the Xsigo Management System, without entering the data center.

In addition to running as a stand-alone system, XMS has been integrated with the VMware Infrastructure Client. This integration enables complete management of virtual machines and virtual I/O from a single user interface.

Traffic Engineering

All I/O modules enable fine-grain Quality of Service (QoS) including Committed Information Rate (CIR) and Peak Information Rate (PIR) settings that are configurable for each virtual resource. The 10Gb Ethernet module also allows QoS parameters, access controls and classifiers to be configured per flow based on the type of traffic.

Fast Server-to-Server Communications

Server-to-server communications benefit from the high-speed and low-latency InfiniBand fabric. By sending information from one vNIC to another, servers move traffic at nearly 5Gb per second, with a latency of less than 30 microseconds, using standard TCP/IP protocols. This performance enhancing feature for Ethernet-based applications requires no modifications to the applications themselves or to the TCP/IP stack.

InfiniBand protocols are supported as well. For high performance computing applications, the Xsigo I/O Director supports MPI with a latency of 3.5 microseconds for 0-byte messages and a bandwidth of 846MB per second for 4KB messages.

High-Availability Ethernet Connectivity

The Xsigo I/O Director supports a layer-2 High-Availability (HA) solution that binds together two vNICs (primary and secondary) through an HA group name and MAC address. The primary vNIC is the active, online vNIC that supports traffic. The secondary vNIC is a live standby which takes over transmitting and receiving traffic if the primary vNIC fails.

The primary application of Xsigo HA is to protect against system level failures by configuring one vNIC on two Xsigo I/O Directors. However, it is also possible to configure HA to protect against the following failures:

Module-level failures: A single vNIC can be configured on two different modules within the same Xsigo I/O Director chassis. If one module fails, the other will take over and support traffic.

Port-level failures: Two vNICs may be configured on the same module. If one connection fails, the other will take over and support traffic.

High Availability Fibre Channel Connectivity

Xsigo virtual I/O supports multi-pathing software such as PowerPath. For high-availability, two virtual HBAs are deployed in each server. Each server is connected to two I/O Directors.

Management Highlights

- Maximum vNICs per server: 32
- Maximum vHBAs per server: 32
- No server re-boot needed when deploying vNICs and vHBAs.
- QoS is configurable per vNIC and vHBA.
- Hardware enforced QoS.
- All I/O is portable among servers.

I/O Director Highlights

- Bandwidth to each server: 10-20Gb/s
- InfiniBand ports: 24
- Max number of servers: No limit
- Total aggregate bandwidth: 780Gb/s
- I/O Module slots: 15
- I/O Module options:
  - 4 x 1Gb Ethernet
  - 1 x 10Gb Ethernet
  - 2 x 4Gb FC
  - 1 x 10Gb InfiniBand
- Serviceable design:
  - Hot swappable and redundant fans, power supplies
  - Hot swappable I/O modules
  - Field replaceable InfiniBand fabric board
  - Passive mid-plane

Server Management

Managing servers connected to a Xsigo I/O Director is nearly identical to managing servers using traditional NICs and HBAs. Virtual NICs each have a permanent MAC address and will be assigned an IP address, either directly or by DHCP. The virtual NIC interface will appear to the operating system the same way that a traditional NIC does.

Virtual HBAs are configured for a Fibre Channel network and storage devices, just as with conventional HBAs.

In either case, interface configurations may be verified using standard commands ("ifconfig" in Linux or "ipconfig" in Windows).

The only Ethernet interface that cannot be virtualized by the Xsigo I/O Director is the IPMI or out-of-band management network. However, a traditional Ethernet management network can be used side-by-side with Xsigo's virtualized devices.

Server Load Balancers or other network appliances will work seamlessly with virtual NICs. They will see the virtual NICs on the servers exactly like traditional NICs. The Xsigo system will be transparent to the network appliances looking from the network toward the server and to the servers looking toward the network.

Xsigo Connectivity

Two options are available for connecting servers to the Xsigo I/O Director:

- **Copper:** InfiniBand copper cables are available in lengths up to 30M. The cables are round and slightly thicker than a standard Cat5 Ethernet cable.
- **Optical:** InfiniBand optical cables are available in lengths up to 100M.

Deploying Virtual I/O

Virtual I/O is easier to deploy than physical I/O. There is less cabling and there are fewer cards, and all configurations are managed from a single console. Wizards guide the deployment of virtual NICs and HBAs.

- **Install cards and cabling:** Each server is configured with one Host Channel Adapter (HCA) card, or two cards for redundancy. These cards are off-the-shelf devices designed for PCI-Express busses. Cards are connected to the I/O Director via an InfiniBand cable.
- **Install driver:** As with any NIC or HBA, virtual I/O is installed on the server via a driver.
- **Connect networks and storage:** Ethernet routers and SAN Directors are connected to the virtual I/O Director via standard Ethernet and FC connections.
- **Deploy virtual NICs and HBAs:** All virtual I/O is managed from a single screen. Wizards permit virtual NICs and HBAs to be deployed on the fly, without server reboots.

Large-Scale Deployments

The Xsigo I/O Director supports large-scale deployments via an expansion switch. With 780Gb of aggregate bandwidth, a single I/O Director can accommodate the I/O needs for hundreds of servers. The Xsigo IS24 Expansion Switch™ provides connectivity for those servers by enabling fan out from the 24 ports found on the I/O Director. Each expansion switch provides 24 additional ports, one or more of which provide a link to the I/O Director. When multiple links are used, the system automatically aggregates them to fully capitalize on the available bandwidth. The remaining expansion switch ports are then connected to servers.

For example, a fully redundant configuration would include:

- 120 servers
- Two I/O Directors
- 12 IS24 Expansion Switches
- Four links from each switch to the I/O Director (40 or 80Gb aggregate bandwidth)

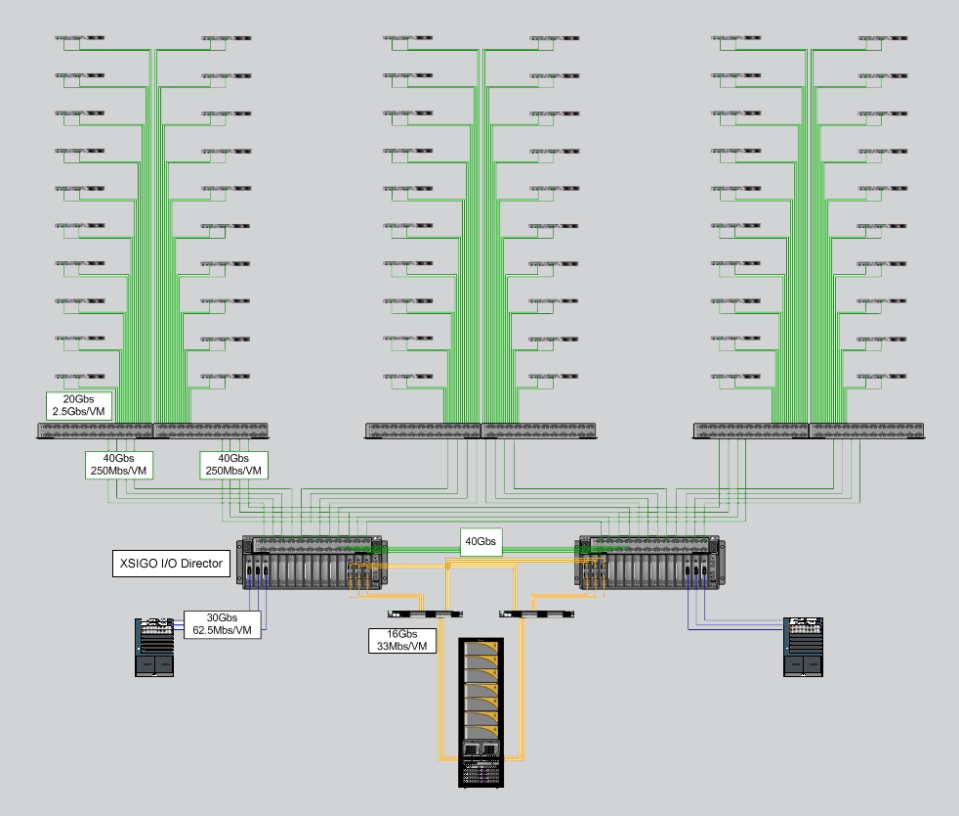

120-server deployment with virtual I/O.

High Availability with Virtual I/O

With virtual I/O, redundant connectivity is achieved exactly as it is with conventional I/O: using redundant devices. Each server is configured with two HCA cards, and two cables, connected to two I/O Directors. Virtual NICs and HBAs can be deployed in HA pairs and configured for automatic failover. Storage software such as PowerPath works with virtual HBAs exactly as it does with physical ones.

The simplicity of establishing HA connections with virtual I/O allows HA to be employed across all connectivity, a significant benefit for application availability.

Cost Savings

I/O virtualization eliminates up to 70% of the components associated with server I/O. This reduction in complexity typically results in these savings:

I/O Capital Costs	35–50% savings vs. traditional I/O for servers that are SAN attached. Savings include 70% fewer I/O cards, cables, FC and Ethernet switch ports.
Server Capital Costs	Savings can exceed $5,000 per server because 1U or 2U server can be deployed rather than 4U servers. Savings include reduced server, I/O, and space costs.
Management Expense	Centralized, remote I/O management means 90% less time is required for routine management tasks.
Power	70% fewer I/O cards and switch ports mean 40 watts can be saved per server, or 80 watts per server if you include power conversion and cooling.

For a 120-server deployment, in three years the savings can amount to over $1,000,000. Furthermore, because capital costs are lower than traditional I/O, the payback in new installations is immediate.

Operating Expenses		**Traditional I/O**	**Virtual I/O**
I/O Management Cost	Per server, per year	$750	$100
Power cost related to I/O @ $0.10 per KW hr	Per server, per year	$140	$70
Total annual expenses per server		$890	$170
Total Expenses			
I/O capital cost	Per server	$4,500	$3,000
Server capital cost (4U server vs 2U)	Per server	$10,000	$5,000
Operating expenses over 3 years	Per server	$2,670	$510
Total 3-year expense		$17,170	$8,510
	Savings per server		**$8,660**
	Savings per 120 servers		**$1,039,200**
	% Savings		**50%**

Summary

Server virtualization imposes new demands on server I/O. Servers now require increased bandwidth, more connections to more networks, and greater management flexibility. I/O virtualization delivers on all of these requirements, at reduced capital cost.

The concept behind server virtualization is simple: Take a single resource and dynamically allocate it to multiple uses. In the process, costs fall and efficiency rises. I/O virtualization does the same for I/O resources. Beyond the cost and efficiency benefits, the resulting environment achieves the same "on demand," highly mobile, easily scalable management model found with server virtualization.

These two technologies, in conjunction with storage virtualization, finally deliver on the promise of the utility data center.

Headquarters:

Xsigo Systems, Inc.
70 West Plumeria Drive
San Jose, CA 95134
USA
Tel: 408-329-5600
Fax: 408-329-5611
www.xsigo.com

Japan:

Xsigo Systems Japan K.K.
Level 7 Wakamatsu Building
3-3-6 Nihonbashi Honcho
Chuo-ku, Tokyo 103-0023
Japan
Tel: +81-3-6202-7484
www.xsigo.co.jp

United Kingdom:

Xsigo Systems, Ltd.
No. 1 Poultry
London, EC2R 8JR
United Kingdom
Tel: +44-0-7917-763156

Germany:

Gartenstrasse 79
D-85354 Freising
Germany
Tel. +49-8161-88-73-94
Fax +49-8161-88-73-85

Australia:

119 Willoughby Road
Crows Nest Sydney, 2065
Australia
Tel: +61-0-2-8587-0050

WP-0108