Gartner Research

Publication Date: 15 May 2009 ID Number: G00167114

Case Study: Cloud Provider Achieves Ten Times Lower Total Cost of Ownership Over FCoE With Input/Output Virtualization

Joe Skorupa

This case study details how a global online transaction service provider leverages server virtualization and external input/output (I/O) virtualization in its colocated data centers to reduce server provisioning time from weeks to hours and to reduce on-site support expertise requirements to "night watchman" levels. When deploying server virtualization, companies should consider I/O virtualization choices, particularly external I/O virtualization, as ways to significantly reduce costs, and improve agility and performance.

Key Findings

- Choosing the optimal I/O virtualization approach can significantly increase the benefits of server virtualization projects. External I/O virtualization can provide significant benefits over server-based solutions.
- External I/O virtualization can allow I/O resources to be pooled, expanded and assigned to individual servers or virtual instances in a more granular and dynamic fashion than traditional hardwired interfaces.
- Companies that focus on total cost of ownership (TCO) and "within the rack" approaches can achieve significant operational advantages over businesses that attempt to roll out a new end-to-end data center architecture. Achieving this leverage may require looking beyond conventional approaches and established vendors.

Recommendations

- Start with a clear set of goals and measurement criteria for your virtualization efforts. Apply these criteria across the project, from initial technology selection through ongoing operation.
- Ensure that your vendors (server, virtualization and network infrastructure) are aligned with your goals and are not just focused on pushing you toward a particular technology offering. Consider agility, lock-in and ongoing support costs as important evaluation criteria of any I/O virtualization approach.
- Don't be seduced by the latest technology hype. Understand the trade-off inherent in your options and don't be afraid to consider creative solutions.

WHAT YOU NEED TO KNOW

External I/O virtualization can significantly reduce TCO, from initial provisioning through ongoing operations. When I/O virtualization is combined with server virtualization, server provisioning times can be reduced by 99% and networking capital costs can be reduced by 50% or more.

CASE STUDY

Introduction

Enterprises are struggling to reduce their operating costs and to increase agility. Often they discover that their approach to data center infrastructure design and implementation conspires against their goals. Many are looking to server and I/O virtualization as a means to overcome the limitations of their current infrastructure and operations but they are discovering that the obvious approaches are immature and not viable. However, there are approaches based on mature technologies such as VMware and InfiniBand (IB)-based I/O virtualization that can be implemented today. The organization profiled in this case study took such an approach and has realized the desired benefits. According to the company's IT architect, "The difference was like night and day. With the old approach, the cabling was a nightmare and server provisioning took weeks. With the VMware and Xsigo approach, time to provision a new server dropped from 216 hours to two hours."

The Challenge

This U.S.-based company provides global online transaction services to the travel industry. It is a division of a Fortune 500 company and processes transactions of more than $10 billion per year. The IT infrastructure is distributed across North America, Europe and Asia, with each data center housing more than 1,000 servers and handling transactions for multiple clients.

The U.S.-based IT architect had an "Aha!" moment when he spent months on site in the U.K. to bring up a new data center. He realized that his company was spending far too much time and money commissioning and managing data centers because each design was custom-made and inflexible. Furthermore, much of the expense was directly tied to on-site installation of equipment, and network and power cabling. Once operational, the data center required skilled on-site personnel to deal with any required repair or change. This realization drove him to develop an approach to data center commissioning and operations that required the colocation facility to provide only floor space, air conditioning and power. Network connections would now be needed only to the virtualization director core. Furthermore, the new model assumed that local technical support would be no more skilled than the typical night watchman.

Approach

This new approach was a key part of an overall data center modernization project with the goals of increasing agility and performance, while reducing power, cooling and floor-space costs. Virtualization and a move to high-end multiprocessor servers enabled two rows of low-end servers to be replaced with two racks of Sun Fire X4600 high-end AMD-based servers running VMware, a suite of custom applications and Microsoft Windows Server software. Each application server is capable of hosting more than 150 application virtual machines (VMs). Larger Microsoft SQL Server databases remain hosted on separate physical servers. The applications are very I/O-intensive and very latency-sensitive, and downtime is unacceptable.

To meet the newly developed cost and provisioning goals, server racks were fully configured before they were shipped to the colocation facility, where they were bolted to the floor and plugged into power and external networking. To cost-effectively support changing client workloads and new business opportunities, a move from one application per server to a fully virtualized CPU model was required. As a result, server I/O had to be able to be reassigned as needed and without recabling the rack. Additional complexity was introduced because virtualization software often requires multiple network interfaces per server. The new approach also had to mitigate this complexity.

After considering alternatives, external I/O virtualization via IB and external network gateways was selected. Mellanox Technologies, the leading provider of IB silicon, was contacted for referrals. Mellanox pointed the company to its system partners, Cisco, Voltaire and Xsigo Systems. The parent corporation is a longtime Cisco customer, so Cisco was selected as the vendor of choice. However, Cisco was unable to supply a suitable solution. As a result, Xsigo was selected.

I/O virtualization isolates logical I/O information (Ethernet Media Access Control [MAC] address, Fibre Channel [FC] World Wide Name) from physical servers and adapters. This allows I/O identities to follow VMs as they are moved from server to server. External I/O virtualization provides the ability to "cable it once from the server to the I/O gateway" and then add external I/O technologies as they mature and as they are required. With this approach, I/O can be added or changed without having to open the servers to add adapters and without the need to recable. External I/O virtualization can also allow I/O resources to be pooled, expanded and assigned to individual servers or virtual instances in a more granular and dynamic fashion than traditional hardwired interfaces.

In the past 12 months, a number of other external I/O virtualization offerings have appeared. Some of these are from emerging companies and some are from more established server and networking vendors, although some work only with servers from a single vendor. External I/O gateways based on Peripheral Component Interconnect (PCI) Express are favored by a number of startups, including VirtenSys and NextIO. While this approach shows promise, products are not yet mature. Server and networking vendors HP, Dell, Cisco, Verari Systems and Blade offer I/O virtualization. HP uses switching technology developed in-house to virtualize FC and Ethernet (but not InfiniBand) connections across multiple racks of HP servers. Cisco uses the not-yet-finalized Convergence Enhanced Ethernet (CEE)/Fibre Channel over Ethernet (FCoE) standards for I/O virtualization for FC and Ethernet (but not InfiniBand) within a rack and across a number of racks of Cisco blade servers. Dell resells the Xsigo solution, and Blade uses its embedded blade and top-of-rack switches to virtualize Ethernet server I/O across an entire data center. Blade also provides its technology to a number of OEM customers, including HP and IBM. All these approaches have software links to VMware to provide some degree of automated provisioning.

Results

The Xsigo product was tested and it performed as advertised. Xsigo was awarded the business. During the testing, IB10 (which operates at 8 Gbps) performance was compared with network interface card (NIC) teaming across four one-Gigabit-Ethernet (GbE) NICs. Table 1 shows a comparison of teamed GbE and IB10.

Table 1. Comparison of Teamed Gigabit Ethernet and InfiniBand 10

I/O Connection	Throughput	CPU Utilization
4x1 GbE	400 Mbps	40%
IB10 (8 Gbps)	7,000 Mbps	2%

I/O Connection	Throughput	CPU Utilization
Note: Newer 10 Gbps network interface cards with enhanced software typically provide better performance than the 4x1 solution, but at higher prices than IB10 adapters. GbE = Gigabit Ethernet Gbps = gigabits per second IB = InfiniBand I/O = input/output Mbps = megabits per second		

Source: Gartner (May 2009)

According to the IT architect, "There are no more patch panels or pulling cables. The Xsigo solution worked so well that we were able to recover from the initial project delay and finish on schedule. Our data center design was transformed from a liability to a competitive advantage. Our salesmen use our ability to rapidly provision servers as a competitive differentiator to win new business."

With the new approach, each server has two IB adapters and is dual-attached to a pair of redundant Xsigo I/O directors that provide local switching, as well as gateways to external Ethernet, IB and FC resources. Storage access is via Network File System (NFS) and Internet Small Computer System Interface (iSCSI). In time, additional network connections, including FCoE, will be available. As the installation grows, the servers will be attached to IB leaf switches that will perform aggregation. The leaf switches will be connected to the pair of Xsigo I/O directors. When the company's storage vendor adds support for Internet Protocol over IB (IPoIB) or NFS/remote direct memory access (RDMA), storage access will migrate to native IB40. Figure 1 shows the before and after server, I/O and cabling configuration.

Figure 1. Before and After Server, Input/Output and Cabling Configuration

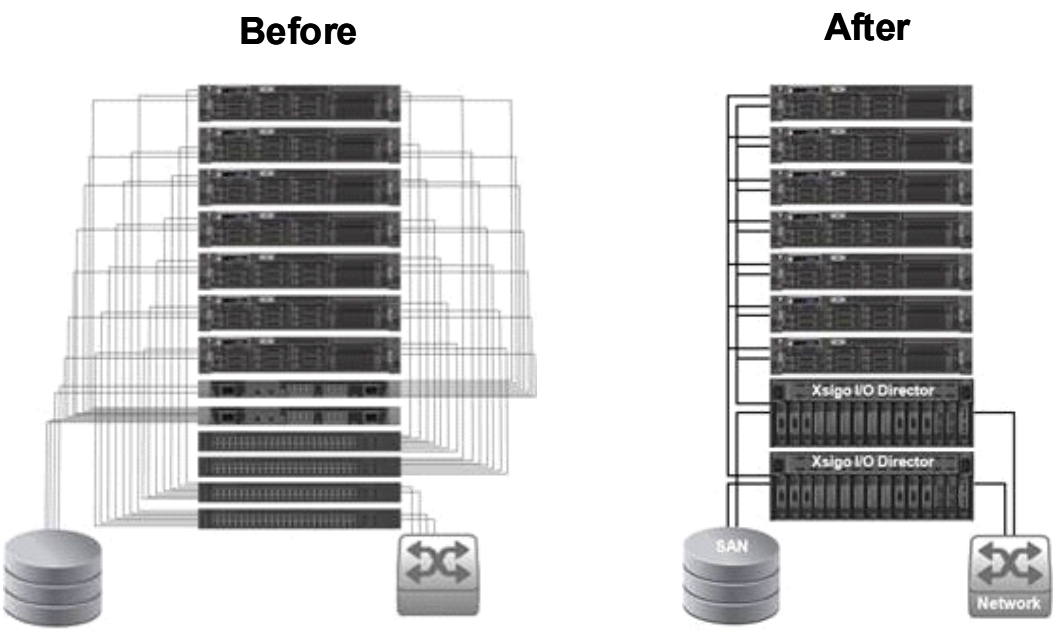

SAN = storage area network

Source: Gartner (May 2009)

Gartner

According to the IT architect, "We don't have to contact the networking team when we want to add a connection to the network. All they have to do is deliver the requested [virtual LANs] VLANs to the 10 GbE switch ports on the Xsigo Director and we do the rest via management software."

Server I/O is defined via the Xsigo management console. Virtual NICs (VNICs) in the servers are mapped to physical Ethernet and FC adapters in the Xsigo Directors. Multiple VMs can share a single physical adapter. Bandwidth and quality of service (QoS) parameters are configurable. If VMware VMotion relocates a workload to a different physical server, the virtual I/O is automatically remapped by software. I/O for new VMs in any data center can be remotely provisioned in minutes. On-site support requirements are limited to the ability to replace failed hardware.

As few workloads can saturate a 1 Gbps Ethernet adapter or 2 Gbps FC host bus adapter (HBA), many VMs can share a single 10 GbE or 8Gb FC HBA in the Xsigo Director. As IB10 adapters are available for approximately $120 and switch ports are around half the price of 10 GbE ports, this approach lowers equipment costs and power and cooling loads.

Traffic between VMs on the same machine is switched by the embedded VMware vSwitch. Traffic that is destined for another physical machine is mapped to the same virtual port on the Xsigo Director and is switched there. The high-speed, low-latency IB connections ensure good performance and prevent bottlenecks. In the worst case scenario (and the least likely one), the traffic is routed between Xsigo Directors across the data center's 10 GbE network.

According to the IT architect, "We are focused on driving the market. With the IB solution *we* get 10 times the performance for a quarter of the price of FCoE."

The database servers employ very large caches to reduce latency and improve performance. The company is also a big proponent of solid-state storage, having already deployed 8TB and expecting to aggressively adopt more in the next two to three years to further reduce latency and improve performance.

According to the IT architect, "Xsigo products have been in production since 2007 and have been extremely reliable and bug-free. We have not had one outage on our dual-switch configuration with VMware. This allows us to write aggressive service-level agreements (SLAs) with penalties; essentially making the customer's business unattractive to our competition. The only notable bug concerned [Simple Network Management Protocol] SNMP [User Datagram Protocol] UDP traffic and Xsigo provided a fix in 48 hours."

Opportunities to further improve the system have been identified. The company is working with Microsoft to extend Winsock Direct to the Mellanox NIC drivers to further improve performance.

The company is also working with Xsigo to implement direct server-to-server networking via the server IB connections and leaf switches, thereby offloading the Directors and further reducing latencies. Xsigo was initially resistant to extract perceived value from the Directors, but has now agreed to add the requested feature to upcoming software.

Beyond the initial difficulties with Cisco's products, the biggest challenges have been related to personnel. When the new model was proposed, many people couldn't conceptualize the new model because of its seemingly radical nature. Some of the staff still do not fully understand the full potential the solution provides to the organization, but are slowly accepting the changes.

Critical Success Factors

- The customer had clear goals in mind for the entire IT modernization project and drove those goals across traditional technology silos.

- The IT architect enjoyed management support that enabled him to drive changes across organizational boundaries and to force the adoption of new technologies from unfamiliar vendors.
- The customer was willing to consider less mainstream solutions to achieve its goals.
- When the initial solution provided by a trusted incumbent vendor failed, the company changed vendors rather than delay the project or adjust the goals to what the vendor could provide.

Lessons Learned

- Commitment to a set of clear, measurable goals and an open mind about vendors and technologies can lead to an "order of magnitude" improvement in price/performance and a sustainable business advantage based on agility, improved customer satisfaction and lower costs.
- It is critical to look across technology and management silos to affect substantive changes in data center architecture and operations. However, this approach can transform IT operations from a liability to a strategic asset.
- Giving the affected team members time to understand and adapt to the new approach increases acceptance and long-term benefits.
- Some team members will never make the transition to the new approach and must be moved into other roles.

RECOMMENDED READING

"I/O Is the New Frontier of x86 Virtualization"

"Emerging Technology Analysis: Fibre Channel Over Ethernet, Networking and Communications"

"The Folly of Fibre Channel Over Ethernet"

REGIONAL HEADQUARTERS

Corporate Headquarters
56 Top Gallant Road
Stamford, CT 06902-7700
U.S.A.
+1 203 964 0096

European Headquarters
Tamesis
The Glanty
Egham
Surrey, TW20 9AW
UNITED KINGDOM
+44 1784 431611

Asia/Pacific Headquarters
Gartner Australasia Pty. Ltd.
Level 9, 141 Walker Street
North Sydney
New South Wales 2060
AUSTRALIA
+61 2 9459 4600

Japan Headquarters
Gartner Japan Ltd.
Aobadai Hills, 6F
7-7, Aobadai, 4-chome
Meguro-ku, Tokyo 153-0042
JAPAN
+81 3 3481 3670

Latin America Headquarters
Gartner do Brazil
Av. das Nações Unidas, 12551
9° andar—World Trade Center
04578-903—São Paulo SP
BRAZIL
+55 11 3443 1509

Gartner