# I/O Virtualization:

## The next step in data center virtualization

A report by Virtualization Evangelist David Davis

Data centers across the world are undergoing a major transformation - they are striving to lower costs, improve operations, and obtain increased agility. Virtualization is the vehicle of to deliver these objectives. Since its inception, virtualization has been the major enabling technology that improves how we consolidate compute resources, deploy servers and improve the scale of our data center applications.

In an era where data center managers and administrators are endlessly tasked with lowering TCO (total cost of ownership) and are measured against the most stringent of service levels -- the need has arrived for an innovative technology that meets the challenge of scaling virtualized applications without bankrupting the business. However most data center administrators agree that lack of sufficient I/O and memory, not CPU is the barrier to achieving higher performance, scale and virtualized server density.

Countless data center administrators spend precious time and valuable resources provisioning physical I/O directly to a single server; often over-provisioned in fear they will not be able to sufficiently support and supply the required I/O to meet expected applications service levels. For example, they may provide each server with its own dedicated 10 Gigabit Ethernet (10 GbE) network interface card, a dedicated network switch port, and cable. This practice is common and causes a “static” 1 to 1 relationship between resource and server.

These types of I/O resources are traditionally PCIe (PCI Express) based adapters and allow servers to connect to high-speed networks and high-performance storage. PCIe is a well-known, well-understood and universally accepted server based technology, used to connect and access resources such as network cards, Fibre Channel host bus adapters (HBAs), RAID (redundant array of independent disk) controllers, and most recently -- PCIe based SSDs (solid state drives) adapters.

### Enter I/O Virtualization

Through use of disruptive technologies, such as hardware-based I/O virtualization (IOV), data center administrators can now pool high value, in-demand I/O resources and share them across multiple servers, rather than a single resource dedicated to each server. This is possible without sacrificing the native capabilities of the I/O resource, which make it valuable to the administrator, such as offloading checksum operations through hardware acceleration or improving how the card buffers information when sharing with resources amongst virtualized servers.

Servers have access to a virtualized instance of a traditional physical I/O resource. Sharing the physical I/O resource means communications between servers connected will achieve higher performance and improved resource utilization which drastically lowers their TCO. Furthermore, centralized provisioning and management also improves the performance and utilization of the shared I/O resource. All servers now have access to I/O resources that previously, for economic and/or logistical reasons was not possible. The performance benefit alone is a huge improvement. For example, block copies within a shared data store, supporting operations within a clustered database or virtual machine live migration between server hosts. This operation can be achieved at local speeds/feeds, as if the storage resource was directly attached and local to the server host itself.

In the past few years new technologies have emerged that enable consolidation in the storage and network layers. Advancements in these areas of the traditional data center have allowed administrators to do more with less and provide much better application services and infrastructure deployment. Their key challenge throughout this period of transformation has been to improve the way they manage the resources they depend upon every day. I/O Virtualization (IOV) provides new ways for administrators to share critical resources; its use extends the data center management paradigm and the subject of this report.

## What is I/O Virtualization?

I/O Virtualization (IOV) is the logical consolidation of I/O adapters or peripherals. IOV allows administrators to remove the physical dependency of a specific I/O adapter from the server. The upper layer protocols are abstracted from the physical media. To the operating system and applications the I/O resource still appears as if physically installed, but in reality it is a virtual representation of the physical resource.

**What does this really mean to you and how does it help you?**

You are probably familiar with server consolidation using virtualization and its benefits: you turn a physical server into a virtual server and consolidate multiple virtualized servers onto fewer hosts. IOV provides many of the same benefits, as the same concept applies. IOV consolidates the storage and network adapters (and associated physical cabling) to a single connection. This results in ease of management, improved agility, and tremendous cost savings.

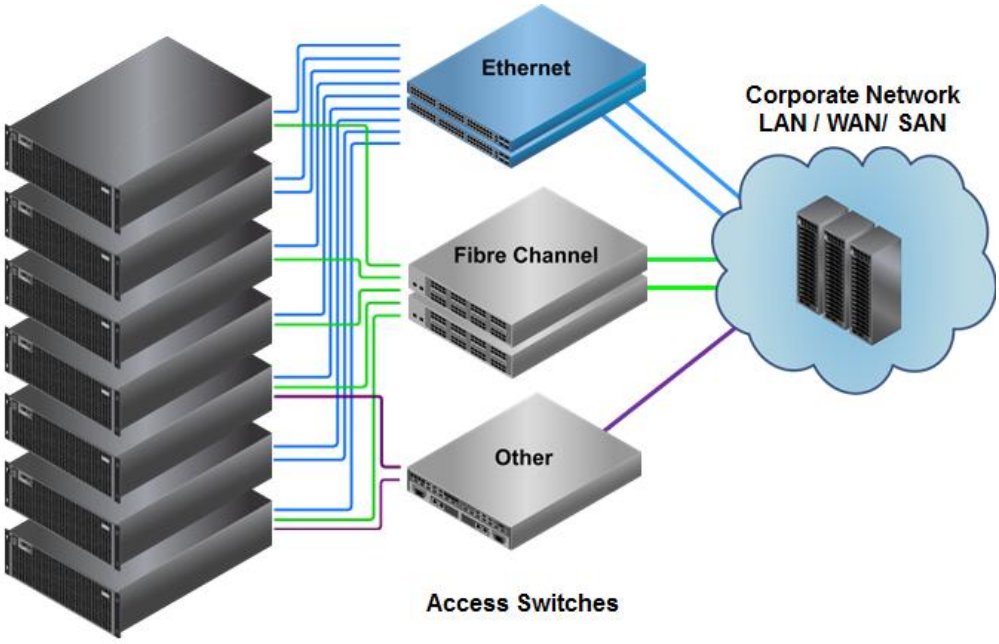


**The traditional physical server I/O infrastructure**

In a traditional configuration a typical physical server may have two to four Ethernet connections (for access to IP based storage, management access, and general purpose application traffic) and at least two FC connections to access storage resources. An infrastructure designer might roll this out as the standard server I/O configuration. There are two immediate problems with this:

- Physical complexity, which results in "hard to manage/troubleshoot" scenarios
- I/O adapters have "1 to 1" relationship to physical servers, resulting in very poor resource utilization

Having a total of four to six cables (simply for I/O traffic) going to a single server is expensive, difficult to troubleshoot, and time consuming to install. Consider also the cost of the multiple network and storage adapters in each server and the six different connections, spread across the core Ethernet and FC switches!

If a critical production server runs out of I/O bandwidth and needs additional capacity, this means shutting the server down, adding a NIC or HBA and a new cable run to the switch. Then powering the server back-up, causing downtime for the end users and lots of work for the administrator. Furthermore, the physical server may be limited in the number of available PCIe slots to add new I/O adapters. Whereas larger servers may have several available slots, with the advent of very dense computing platforms (1/2U servers), the administrator may only have a single available slot.

With I/O virtualization implemented the benefits are obvious:

- Drastic reduction in the amount of physical cabling
- Simplified troubleshooting of server connectivity
- Cost savings in cable, NICs, HBAs, and Ethernet/FC switch ports
- Ability to dynamically add and remove I/O connections to servers

In other words, no matter which way you look at it, IOV eases the management burden, saves money, and gives the business the flexibility (and greater uptime) that it needs.

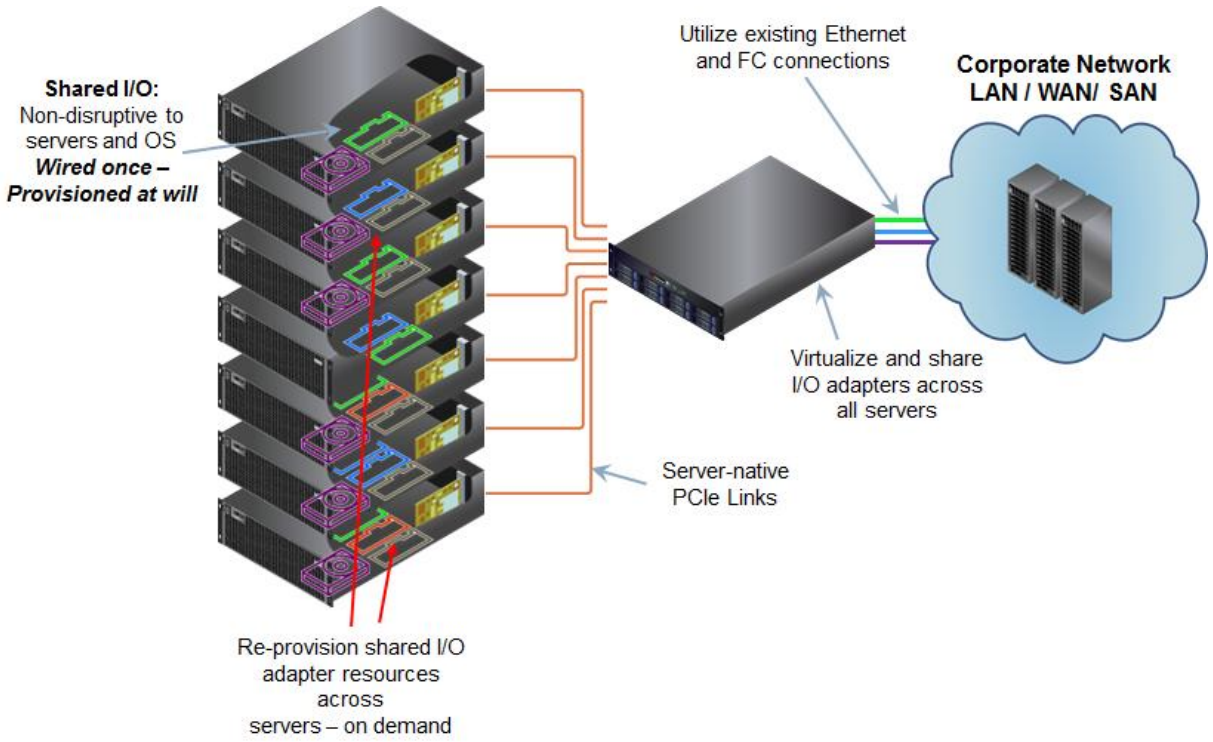


**A physical server I/O infrastructure with IOV implemented**

## Server Virtualization - Exasperating the Problem

A recent survey by a leading analyst states that "78% of servers will be virtualized by end of 2011". VMware® themselves state that roughly 50% or greater of global businesses leverage their product to virtualize servers and application workloads now. When implementing server consolidation through virtualization, the goal is to run more virtual machines on fewer physical servers. A big part of this goal is to drive up hardware utilization on the physical servers to achieve the highest ROI. However, even if you have enough capacity for your company's virtual machines from day one, an ever-expanding number of virtual machines and dynamically changing load can quickly cause I/O bottlenecks and performance issues.

I/O Virtualization gives a virtual infrastructure (and the data center administrator) what is needed - flexibility to add and remove I/O connections as the demand on the virtual infrastructure changes. This flexibility enables the infrastructure to become more agile.

At VMUGs (VMware User Group meetings) across the country VMware representatives are attempting to address the hard to virtualize server applications (Microsoft Exchange, Oracle, SAP). CTOs and heads of IT department desire to move all server applications to virtualization, not only for the TCO benefits, but for the agility and redundancy rapid deployment they provide. There is a growing move to utilize server virtualization even in a 1:1 P2V (physical to virtual) deployment. The IT administrators, however, are reluctant, due to the inability to guarantee the adequate I/O resources and bandwidth required.

IOV offers intelligent management and configuration of all I/O connections and prevents I/O bottlenecks, enabling administrators to dynamically add and remove traditional I/O connections.

## Approaches to minimizing I/O management with IOV

Just as with server virtualization, there are multiple approaches to IOV. However the goal of all IOV solutions is to reduce I/O infrastructure and increase business agility; although solutions differ on how they achieve this. A key part of the difference centers around how servers connect to the I/O gateway and how those gateways are interconnected. Let's look at the four solutions in which I/O virtualization is offered today:

### Option 1 – The Encapsulated Network Approach

The encapsulated network approach offers a way to reduce physical cabling and minimizing the total number of I/O adapters per physical server. One network type is encapsulated within another network medium. The network used for the encapsulation may be one of your existing networks or it may be an entirely new network technology that you have never used before.

Primary vendors of the "IOV encapsulated network" solution use either PCIe over Ethernet with 802.1Q Ethernet frames or Infiniband technology to connect the physical servers to the I/O they need. Both of these solutions work to encapsulate I/O traffic inside their medium; this is analogous to how iSCSI traffic is encapsulated inside TCP/IP sessions over Ethernet networks.

PCIe is an extremely well known and trusted server based technology. In the encapsulated network approach, this solution uses standard Ethernet cabling and switches but encapsulates PCIe communications within the Ethernet frame, requiring specialized adapters on the physical server.

Encapsulating I/O using Infiniband requires introducing another I/O technology into the data center. Infiniband cards and cabling are expensive and a less common network technology in virtual data centers. This solution may also require installing a driver on to each of your virtualization hosts - drivers that are not necessarily native to the OS or hypervisor and may be proprietary solutions from the Infiniband vendor.

In summary, the IOV encapsulated network approach has both good and bad points - you may have to use specialized Ethernet or Infiniband cards in each of your servers and also install drivers to use it. Additionally, depending on the network chosen, the IOV encapsulated network approach may introduce new and unfamiliar network switches, adding more cost and require additional personnel to support.

## Option 2 - Reconfigurable I/O

Reconfigurable I/O allows administrators to assign an I/O resource from a shared pool of resources. For example, if a server required a 10 GbE network to perform a high bandwidth related task (such as a file transfer or nightly backup), you could temporarily assign it to a given server and once the task is completed, release back to the pool for another server to use.

Instead of fixed I/O resources per server that require dedicated management, a single pool of resources is managed through one "pane of glass". The resources are dynamic, and can be assigned and reassigned on the fly to any server. Server installs become hours instead of days. Server failovers become minutes rather than hours.

In this approach the I/O resources are not shared but mapped "1 to 1" with the server resource. The ability to reconfigure the I/O resource dynamically makes this an interesting option. For instance, not all servers may need access to a SCSI controller, but perhaps there is a need to back a server up with a legacy backup target technology like a high speed tape drive. You could assign access to a given server, allow the backup process to complete and then move the I/O resource to the next server in line. These processes can be scripted and even automated for additional savings.

In addition, a PCIe based IOV solution (transparent to the host) preserves the functionality of the I/O resource and server platform and virtualization technologies.

## Option 3 - Converged Ethernet / Fibre Channel over Ethernet (FCoE)

Large network vendors have pushed the convergence of FC I/O onto Ethernet networks. Although not exactly considered IOV, the goal of FCoE is to reduce the complexity of the I/O infrastructure. However this also introduces additional administrative and configuration complexity. So in some cases IOV and FCoE are confused, but for the sake of comparison, the report will concede that the two solve similar challenges but go about solving them in different ways.

While converged Ethernet may seem like a good idea at first glance, you will find that, while it has been around for many years, the adoption of FCoE has been slow. Instead of adopting FCoE, most datacenters have moved to iSCSI or have preferred to leave their I/O infrastructure as it is. FCoE hardware is expensive.

## Option 4 – PCIe sharing

The fourth approach to I/O virtualization is to naturally extend the hardware through a native extension of the server PCI bus. This allows sharing of I/O adapter resources through hardware virtualization. The PCIe sharing solution requires a PCIe extension card to be installed in each physical server and a PCIe extension cable connected the server back to the appliance.

Time to implement PCIe sharing solution is minimal and requires no additional drivers to be installed on the physical servers. Because the solution leverages the native drivers for the virtualized I/O resource (for example Intel's x520-SR2 10 Gigabit Ethernet controller or QLogic's QLE2562 8 Gbps Fibre Channel Host Bus adapter) these drivers are already within the operating system or hypervisor. Virtualization hosts (ESX/ESXi servers) will see 10 GbE Ethernet and FC HBA interfaces upon reboot and you can start using these connections immediately.

The PCIe sharing approach, such as that offered by Virtensys, is a logical extension of the server's native capability to use PCIe based resources. It

offers the least disruption and requires the least amount of change to the existing data center environment. It is also the fastest "time to market" of all IOV solutions.

In Virtensys' case, IOV is a use case of their PCIe sharing technology; it provides shared access to traditional I/O adapter resources. Their IOV solution improves the way PCIe adapters are leveraged in the virtual data center.

## Trends in the data center

How can IOV help the average server administrator in terms of easing administrative burden and lowering the cost of deploying new servers (or even reducing the cost of existing servers)? To find out, I conducted a random survey of server administrators. The majority of those surveyed manage a large number of rack-mounted servers and most are already using server virtualization. The results revealed some very interesting trends.

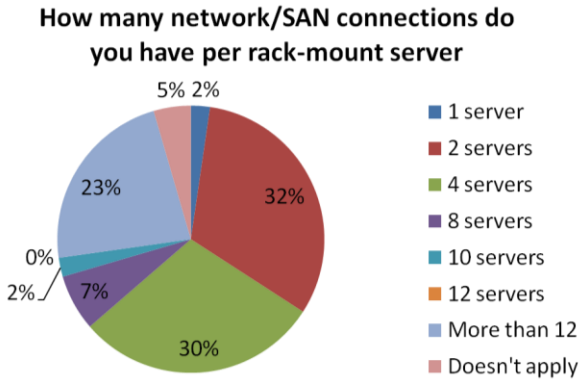


On each of their physical servers, most administrators have more than four I/O connections going to the network and the SAN (20% have as many as 12). On average, each of those I/O connections costs about $1500. Even if we take the lower number of four I/O connections, we find that $6000 is typically spent to connect a single physical server to the network and storage. For the 20% of server administrators who said that they had 12+ I/O connections on each server, this costs about $18,000 per server!

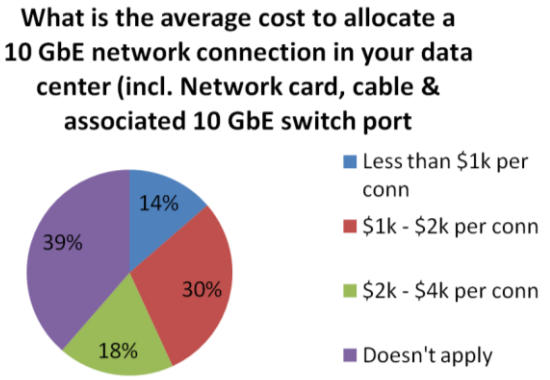


Adopting I/O Virtualization drastically lowers the overall cost of connecting servers to network and storage I/O. IOV also drastically reduces time troubleshooting I/O connections because you have a simple I/O connectivity model.

Regarding cost, as data center planners and administrators seek to lower costs and leverage "commoditized" server, storage, and network resources, technology that is essentially "agnostic" to the environment and allows change of infrastructure is an enabling factor. Cloud computing, both internal and external to the data center depends on the ability to leverage commoditized hardware. Having to deploy a specific server and not having the choice to change to a new make/model or chose an entirely different server manufacturer would be a huge inhibitor. Technology that enables "change without disruption" will, in turn, enable cloud computing adoption.

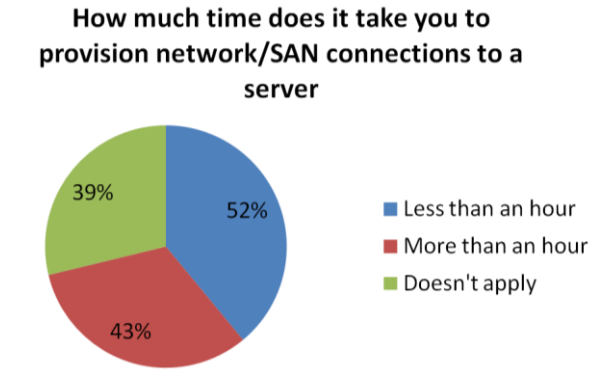


Server administrators also say that, on average, it takes about an hour to provision a new I/O connection to a server and they need assistance from other IT groups such as the Network or Storage teams, to facilitate this.

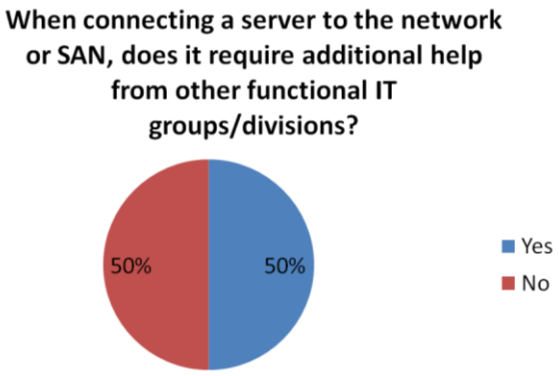


By using IOV, you are able to add new I/O connections to servers in minutes without assistance from another team.

When it comes to reducing downtime, half of server administrators said that their users would suffer downtime when they had to add a new network or SAN connection to a server. Once IOV is up and running, downtime can be eliminated as new I/O connections can be added in minutes.

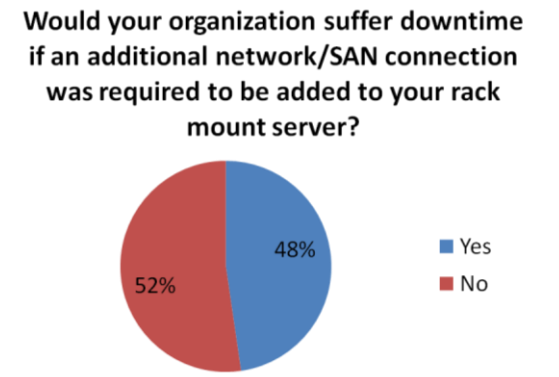


Finally, the server administrators polled told us that, when it comes to I/O, the piece they struggle with the most is physical server cabling. Rack cabling is a huge pain to install and to troubleshoot. In companies that are growing and changing, physical server rack cabling is in a constant state of change and, in many cases, unsightly and disorganized.

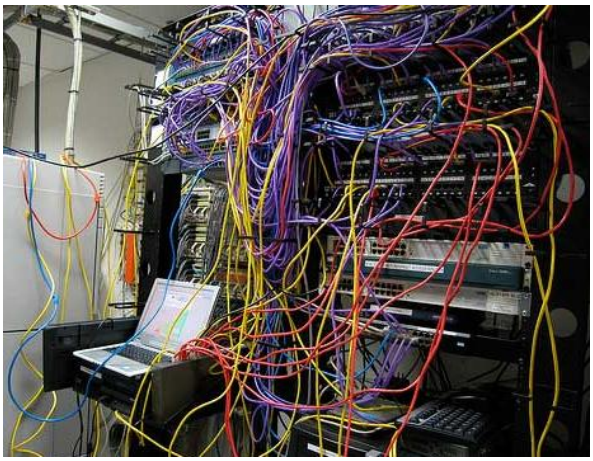

As much as 80% of the physical cabling that runs to servers today can be eliminated with IOV. How much time and expense would be saved by reducing such a huge percentage of your datacenter cabling?

In summary, server administrators spend too much time working with physical server cable problems and companies spend too much money on cabling and suffer from too much downtime as a result of cable problems and time-consuming troubleshooting. I/O Virtualization is the solution!

## I/O Virtualization in action: A Customer Case Study

What about companies who are already using IOV today? One such company is the UK's leading outdoor clothing supplier, Regatta. Regatta recently implemented Virtensys' PCIe sharing solution to deliver IOV to its physical servers running VMware vSphere and Red Hat Enterprise Linux. Regatta's decision to use IOV was based on their desire to achieve the highest virtual machine to physical server consolidation ratio.

Rather than accepting the inflexible, unmanageable, and non-scalable approach of forever deploying additional 1 GbE network interfaces as "Band-Aid" solutions to I/O provisioning, Chris Bulmer, Regatta's Director of IT Operations, was seeking to implement a more robust and agile solution that would allow them to scale applications and critical services beyond the norm. By consolidating and virtualizing their network and storage I/O, Bulmer reasoned, IT would be able to:

- Provision I/O dynamically and remotely in the same manner as a virtual machine
- Make servers "agnostic" to the storage/network resources
- Improve resource utilization of essential devices, achieving even greater VM density (virtual to physical)
- Free server and network administrators from resource intensive "routine" tasks and allow the IT staff focus on greater improvements in data center efficiency from their physical servers.

Regatta's datacenter has migrated from an inflexible collection of 1GbE and 4GB FC HBAs across hundreds of servers to an elegant, flexible and efficient I/O solution that provides them drastically higher VM consolidation ratios.

Since implementing Virtensys' VIO-4000 series appliance, Regatta's IT administrators can now:

- Simply "wire" their servers once and deploy 10 GbE network connectivity "at will"
- Seamlessly add robust/resilient network interfaces to critical applications and services that demand more bandwidth and throughput
- Deliver higher VM densities per host, by delivering more dedicated I/O per host
- Eliminate the physical cabling that had previously been required to meet addressed network connectivity needs
- Leverage simultaneous access to different I/O resources, allowing them to have the flexibility in the future to leverage both 10 GbE and Fibre Channel simultaneously to assist with storage migrations and storage related operations
- Maximize their I/O diversity per server host, yet within the smallest possible physical footprint

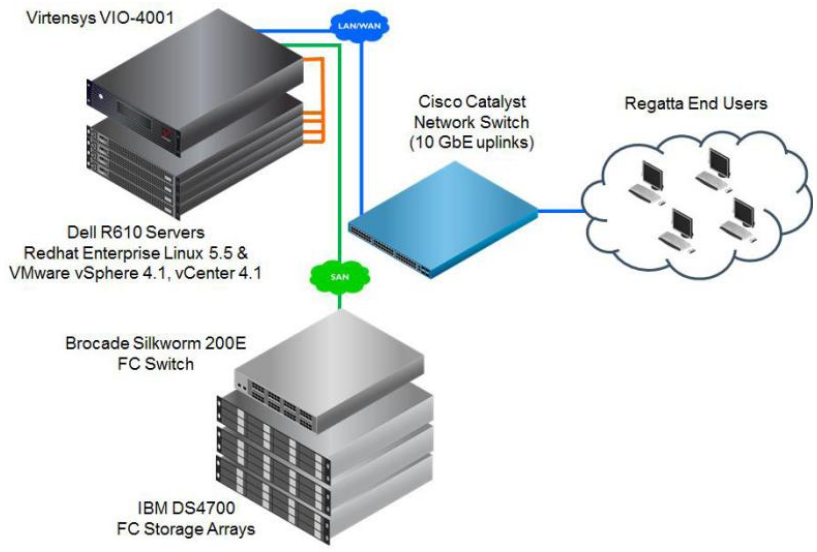

Seeing what IOV is able to do for Regatta is eye-opening. The flexibility, cost saving, and increased VM density is really impressive. As a server administrator myself, I would be willing to bet that the stress level of their server administrators has been reduced, and they are able to now focus on more important and higher profile projects that have a positive return for the company (rather than just maintaining the cable plant).

I recently had the opportunity to sit down and interview Chris Bulmer, VP of IT Operations for Regatta. During the interview I asked some specific questions regarding the goals for his data center and how Virtensys' PCIe sharing solution, played a role in his virtualization strategy:

David Davis (DD):

Chris, can you tell us about your role, your team, data center operation (# of hosts, VMs, and # of administrators) at Regatta?

Chris Bulmer (CB):

I am the VP of IT Operations and am responsible for the technology department at a large clothing manufacturer in the UK that specializes in outdoor apparel. As you can imagine, my team and I are very busy, as our core business isn't IT, so there aren't a lot of us, but there are a lot of IT centric projects -- so we are in-demand often.

Within our datacenter we've got in the region of 60 physical servers in total. Not all of our physical servers run a virtualization layer; the total number of VMs is possibly nearing 100 by now. We're still going through a growth and refresh phase here, so we plan to address the balance and ratios as things progress.

As already mentioned, my team is quite lean but each and every team member is exceptional at what they do, and many are multi skilled. One of the most acclaimed achievements within Regatta is the low turnover of staff. In part this is due to the culture, equally though, because of our adoption of new technology and the innovation that gestates as a result.

DD: Do you have any rough estimate as to how much money was saved by using IOV?

CB: Quantifying this sort of saving at the macro level is straightforward; consider the cost of adding in a new 10G core as we would have had to. Add to the cost of a new core putting the entire high-end 10G cards in each host, this quickly becomes a financial nightmare. We're a well-run profitable company but even so, going to the board and asking for tens of thousands of pounds, ignoring all the disruption etc. is a very difficult position to be in.

On a micro scale we're still working through the overall saving, certainly the amount of tickets we have to resolve due to slow or poor network access has reduced. I can't go into numbers but take it from me the cost savings from operational efficiency are more than enough justification for putting in the Virtensys VIO-4000 series appliance.

DD: Was there a reduction in the number of switch ports and/or FC ports which resulted in cost savings?

CB: Without doubt we had a pain point which was core bandwidth in any case, adding more connections into the core meant replacing the core switches. This is no longer a worry as the bulk of everything causing the concerns is confined to the appliance. On the SAN side we had less of a port constraint, with the appliance we get an 8:1 difference in the ports needed on the SAN.

DD: What is your new VM to host consolidation ratio with IOV implemented?

CB: We're not one of your super aggressive guys with VM to host ratio, we're at about 5:1 right now. The comforting position is we can now scale this up without choking the network; it's simply a case of adding more memory or swapping the servers out for newer generations.

DD: What was the implementation of IOV like?

CB: Truthfully, the implementation was very smooth. Once we'd planned out how and when to move the infrastructure around. There was some preparation work for us to do internally, which, to be fair, we ought to have done in any case.

The deployment of the VIO-4000 series appliance was carried out by the Virtensys team, in less than 30 minutes. They made sure that, once their part was complete, everything on the server and network side was running correctly.

DD: Are all servers using IOV? Do the servers using IOV have a special use case?

CB: Yeah, we have a bit of a mixture in our data center; a number of the newer servers are now on the IOV appliance. We have a very split system connected to it some applications run in dedicated VMs whilst others are on the tin running flavors of Red Hat Linux. We write a lot of our own internal applications simply because of the dynamics of our business and operation.

DD: How long does it take to provision a new Ethernet or FC connection for a server, now that IOV is in place? And how long did it take before?

CB: We can now have a new 10 GbE port available in less than 60 seconds, even my guys that weren't trained on the system can create resources as they need! The Virtensys system is very quick and intuitive to use.

Prior to having the system in it would have taken 30-60 minutes (sometimes more) to simply add a card plus cable the server up. Thankfully we don't even worry about this now with what we've built around the Virtensys appliance.

DD: Did you consider other IOV companies and why did you choose Virtensys over them?

CB: As I'm sure you know there are a number of vendors out there who offer this sort of technology, majority of which are software widgets running inside of a VM. A number of hardware vendors that do have offerings are either too expensive or expected us to change our entire infrastructure to make their particular product fit. For us the VIO was the best all round solution considering cost, port count and performance. It also makes for an easy decision when we have such a level of access to Virtensys' engineering teams.

DD: Have your end users seen any direct benefits from the IOV implementation? And/or do you have any measurements showing increased throughput after the IOV implementation?

CB: As I alluded to earlier, we had a network load issue prior to the installation of the appliance. Our internal stock management system is real-time and pan European across all of the sales teams. This caused problems with transactions on the network between our front-end web servers and the back-end databases. With the adoption of IOV we keep a good bulk of the traffic (roughly 50%) confined to the internals of the appliance. This means the load on our network was literally halved when we switched the appliance on!

The end users see this through less of a need to hit my NOC team all the time, so I guess one hidden benefit is the overall reduction on staff stress levels…!

DD: Did the administrators have a learning curve when implementing IOV?

CB: Not really, we had the skeptics as you might naturally expect, a radically new way of doing things is always going to cause waves in one way or another. However, the box is very intuitive to administer and in very little time you can be creating exactly what you need to get things running. The learning was more in the sense of mindset verses actually having to learn something new. The VIO appliance is truly a fit and forget unit, we only ever touch it to either lookup MAC assignments or to move resource around.

## The Virtensys Perspective

Traditional data centers take a conservative approach when it comes to provisioning and managing I/O connectivity to rack-mount based servers. This typically consists of installing multiple network interface cards (NICs) and storage host bus adapters (HBAs). To keep pace with business demands for greater connectivity bandwidth requires installing more and more I/O resources. This impairs the IT organization's ability to remain agile and responsive to the business needs as time, cost and resource needs for setup/teardown simply becomes greater.

Virtensys' approach to IOV makes I/O provisioning and management simple and transparent. This means that servers can leverage virtual instances of traditional NICs and HBAs, which appear to the server as locally installed physical I/O adapters. I/O provisioning to the servers is managed remotely and done "on-demand". A core value of the Virtensys' appliance is to simplify the I/O provisioning process, thereby allowing IT to address greater operational and capital expense improvements.

Virtensys' award winning VIO-4000 series architecture leverages industry standard PCIe (PCI Express) technology. The PCIe bus is the primary motherboard-level interconnect, connecting the host system processor with both integrated peripherals and add-on expansion cards. It's the native communication bus of all x86 modern systems. Virtensys extends the PCIe bus out of the server with a low power, low profile, and low cost PCIe extender card. Connecting each server to the Virtensys VIO-4000 series appliance with industry standard PCI extension cable (the lengths range typically 2, 3, 5, or 7 meters). The VIO-4000 series hardware virtualizes industry leading high-end PCIe I/O cards (Intel's x520 SR dual port 10GbE Server Adapter, QLogic's QLE2562 dual port 8G FC HBA, and LSI's 9260-4i 6Gb/s SAS/SATA RAID controllers).

## This approach provides a number of obvious advantages:

- The VIO-4000 series presents virtualized I/O and is transparent to the servers and the upstream network, meaning: the vNICs, vHBAs and vDAS appear as if locally installed in each of the servers because they are locally installed. Extending the PCIe bus of each server makes the VIO an intricate layer of all connected servers. For instance, the VMware vSphere host server displays the Intel, QLogic and/or LSI adapters within its configuration. There are no additional drivers to install; rather it uses the drivers already available within the OS or Hypervisor. There is no new technology to learn. Server, Network and Storage administrators know well how to manage Intel NICs, QLogic HBAs and LSI DAS within servers.
- A single physical cable out of the back of each server eliminates the spaghetti factory cabling nightmare in the datacenter. This means that you can simple "wire once and provision at will". This model provides a centralized location of I/O resources eliminates routine task of cabling, firmware updates, and taking down servers to add more bandwidth. Now in seconds, additional FC HBAs, 10GBE bandwidth and storage can be dynamically provisioned.

When I first saw the VIO-4000 series appliance, my first thoughts were, 'how primitive,' until I considered the upgrade potential. Virtensys can theoretically virtualize *any* PCIe device. For instance, their flagship model, the VIO- 4001 (FC and 10GbE) initially shipped with the 2 x Intel dual port 10GbE (40 Gbps of Ethernet) and 2 x QLogic 2462 dual port 4G FC (16 Gbps of FC). It now supports the QLE2562 dual port 8G FC HBA, so within minutes, 16 servers can be upgraded from 16 Gbps to 32 Gbps of shared FC bandwidth. As 40GbE network adapters come to market, in the not-too-distant future; Virtensys will again double the potential throughput and physical port count available on any already purchased VIO appliance.

In addition, a PCIe sharing based IOV solution (which is completely transparent to the host) preserves and, in some instances, amplifies advanced Intel® platform and virtualization technologies. The list of these advantages goes way beyond the scope of this article, but to mention just a few:

- Intel I/O Acceleration Technology is a set of technologies that contributes to increased IO performance.
- Intel QuickData Technology enables data copy by the chipset instead of the CPU, to move data more efficiently through the server and provide fast, scalable, and reliable throughput
- Direct Cache Access (DCA) allows a capable I/O device, such as a network controller, to place data directly into CPU cache, reducing cache misses and improving application response times
- Receive Side Coalescing (RSC) aggregates packets from the same TCP/IP flow into one larger packet, reducing per-packet processing costs for faster TCP/IP processing
- Low Latency Interrupts tune interrupt interval times depending on the latency sensitivity of the data, using criteria such as port number or packet size, for higher processing efficiency
- Virtual Machine Device Queues (VMDq) optimizes the processing of VM data traffic to improve CPU utilization and bandwidth.

As the number of VMs on a server increases, so does the amount and complexity of traffic. VMDq manages the VMs' data traffic efficiently in order to reduce the I/O bottleneck in the system:

- Throughput: Provides an alternative to VMM-based packet sorting, to ease throughput limitations
- Scalability: Creates parallel data I/O paths in the network I/O silicon to avoid performance degradation as the number of VMs increases
- Capacity: Liberates CPU cycles otherwise consumed by packet sorting, making them available to applications

These advances increase server-consolidation ratios, adding to the cost savings associated with virtualization solutions. Put simply, there is a built-in layer 2 Ethernet switch within the Intel x520 10GbE silicon. VMDq allows virtual machines within the same physical host to communicate at low latency - faster than Ethernet line rate speeds.

## Boiling it down…

**Virtensys = VMDQ on steroids:** Virtensys' MR-IOV (multi-root I/O Virtualization) technology hardware virtualizes the same Intel x520 10GbE silicon and shares it among multiple servers. This uniquely allows the VM on server one to communicate directly with the VM on server eight at low latency – faster than Ethernet line rate speeds (9.6Gbps) through a single 10GbE port.

The preservation and enhancement of the above mentioned platform and virtualization technologies creates a unique low-cost cluster server interconnectivity solution. In x86 server architecture, the VIO appliance can be thought of as the "South Bridge" of each server.

More important still -- imagine if all servers in a cluster shared a single South Bridge with 8 x 10GbE ports, Direct Cache Access, Low Latency Interrupt and Receive Side Coalescing capabilities. This creates a powerful cluster server node to node interconnect at PCIe line rate speeds. Utilizing what was designed to enhance data acceleration across the single computing platform; but now optimized across a cluster of interconnected servers providing the ability for greater amounts of bandwidth/IO for 'hard to virtualize' I/O constrained applications.

Finally, because all I/O is managed centrally within the VIO-4000 series, the I/O profile now resides in a single centralized location. Similar to the vMotion of a single VM. A server recovery is a simple matter of moving the I/O profile of one server to another physical box. If a physical server goes down, the replacement server inherits the same I/O profile assigned to slot on the VIO appliance and can be up and running in minutes. Site disaster recovery is a matter of saving the configuration file (a single xml file) of one VIO and restoring onto another VIO. The ability to move a whole rack of servers from one physical location to another becomes a "drop and drag" function of the upcoming Virtensys management. Things are getting *very cool*, *very dynamic*, and soon, *even automated*!

For further information on IOV and PCIe sharing visit: www.virtensys.com or email: info_request@virtensys.com

## Author's Perspective

I/O Virtualization (IOV) is a type of virtualization that server administrators, IT Managers, and CIOs alike need to investigate to learn its real opportunity to offer so many benefits at so many levels:

- Administrators can save time bringing new servers online and troubleshooting existing servers.
- End users can save on downtime when new capacity needs to be added.
- CIOs and the company overall can save money on network and SAN equipment and give the company scalability.

I/O Virtualization has reached a critical maturity level where it is now applicable to the majority of datacenters and applications.

While IOV may still be a relatively new technology and a different approach within the data center infrastructure, it is gaining popularity. And it does this because it solves *real business problems*.

The biggest challenge to adopting IOV technology is IT's reliance on the "business as usual" approach. If IT administrators are looking control OPEX (operation expense), reduce CAPEX (capital expense), and remain competitive, they need to be open to a new way of performing critical tasks. As with all technology adoption, disruptive technologies like PCIe sharing change the way people approach a traditional task.

Adopting a new technology or approach, shouldn't impact the stability of the environment or require vast changes. The best advice to those seeking solutions to solve I/O management issues is to do the research of the various approaches to IOV. Be sure to look carefully at the implementation choices/challenges of the various technology options.

Sharing in the data center has been proven technology approach. It is a better way to achieve scale, resource utilization and even improve performance. As we expand beyond what we know can be shared (server, network, and storage) to other areas of the data center, we realize that it is prime time that we also leverage sharing for PCIe based resources.

Besides just learning about IOV through research, smart companies, such as Regatta, are taking time to gain hands-on IOV experience and find out for themselves what IOV can offer in terms if benefits and cost savings.

**David Davis** is the author of the best-selling VMware vSphere video training library from Train Signal. He has written hundreds of virtualization articles on the Web, is a vExpert, VCP, VCAP-DCA, and CCIE #9369 with more than 18 years of enterprise IT experience. His personal Website is VMwareVideos.com



 