


Tackling the Physical Limitations of the Server – PCIe Sharing

The challenge in the data center today

All too often in today's data centers, administrators are running out of server resources: CPU, memory or I/O (storage & network) resources are all in high demand. A common problem is that the server has no ability to upgrade or move beyond its physical constraints. Smart data center administrators plan the best they can; buying servers that will meet the performance/scale demands of their applications. However, with the rapid growth in server virtualization, the traditional approach to planning your physical data center needs have changed dramatically. Now administrators are sizing servers not just for a single application, rather for many applications, running concurrently with a variety of resource needs that range in complexity.

Server virtualization platforms such as VMware vSphere® have taken the lead in dynamic virtual server resource management for the data center. However there still remains the challenge of improving resource management at the physical layer.

An acknowledged goal for the next generation data center (based on a heavier use of server virtualization and the promise of “hyper-consolidation”) is to minimize the overall physical infrastructure. From a networking perspective, this means reducing the number of switches per row; from a storage perspective, consolidating and reducing physical storage resources to the minimum, leveraging technologies such as compression, de-duplication, and high-density hard drives.



Blade server architecture is a successful and proven solution to help achieve this. Another solution, and an administrator favorite, is the rack mount server. Many administrators are deploying smaller (1U and 2U) physical server footprints to achieve a reduced cost (for example: power, cooling, physical space) within their densely populated rack. Especially as many customers and providers migrate towards a more commodity server driven model (vis-a-vis Cloud computing and/or Infrastructure as a Service – IaaS) this makes better business sense.

However this also raises some questions. If you plan to reduce the physical size of the server:

- What happens to the resource needs for your application?
- What if the server requires access to a high-speed network, using 10 Gigabit Ethernet (10GbE) and a high-performance SAN using an 8 Gbps Fibre Channel (FC) HBA?
- What if you also require access to the highest performance storage in the market today – PCIe based SSDs (solid-state drives)?

The question isn't why would you want all of those capabilities in the same server, but rather how will you fit them inside a densely populated 1U physical server footprint?



Enter PCIe Sharing

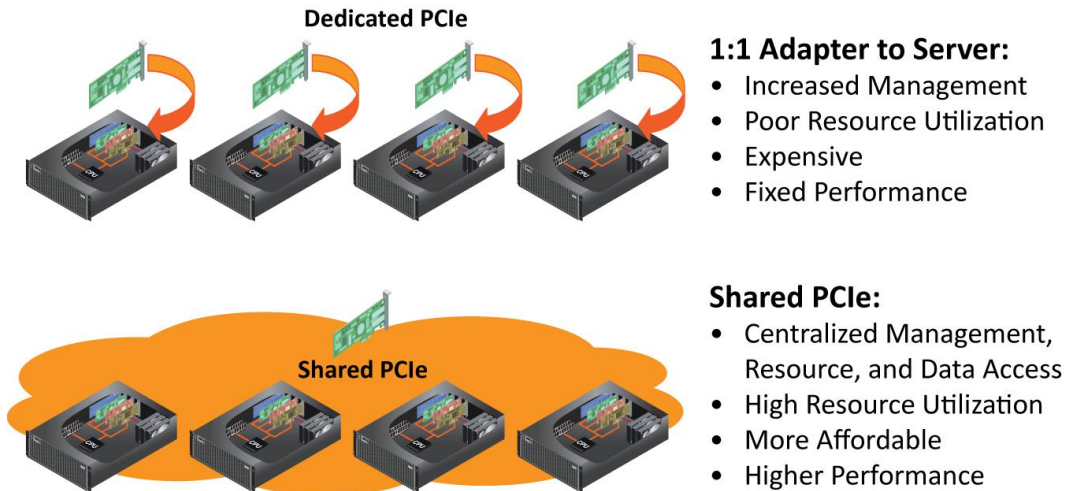
The constraints are obvious. Smaller physical server hosts are “PCIe” slot constrained. Adding more slots to the motherboard is not an option. This means you will not be able to have all the interfaces available simultaneously. Furthermore if you require redundancy (multiple adapters for higher levels of availability), you are completely out of luck.

These types of server I/O peripherals are PCIe (PCI Express) based and allow servers to connect to high-speed networks and high performance storage. PCIe is an industry standard, well-known, well-understood and universally accepted server-based technology. Implementing a PCIe sharing solution that virtualizes multiple physical PCIe I/O adapter peripherals (10 GbE NICs, 8 Gbps FC HBA, RAID controllers and PCIe based SSDs) and shares them across multiple servers overcomes these constraints.

Virtensys PCIe sharing appliance is a proven solution that allows the sharing of physical I/O adapter resources. These types of I/O adapter resources are typically high cost, high demand I/O resources. Sharing them across multiple servers makes sense when you want to lower the TCO and improve resource utilization.

Virtensys' PCIe sharing solution makes this possible without sacrificing the native capabilities of the I/O resource, retaining its initial value to the administrator:

- Offloading compute intensive checksum operations through hardware acceleration.
- Improving how the adapters buffer information when sharing resources amongst virtualized servers.



Sharing provides the **capacity** and **performance** a server/application needs **when** it needs it

How PCIe sharing works

Utilizing Virtensys' PCIe sharing technology, servers gain access to a virtualized instance of a traditional physical I/O resource (such as Micron Technology's P320h PCIe SSD adapter). Sharing the physical I/O resource means that communications between connected servers can achieve higher performance, improved resource utilization and drastically lowers the total cost of ownership (TCO).

Another advantage of the shared resource model is the possibility of centralized provisioning and management. Each server now has access to a range of I/O resources – previously not possible, which has dramatic economic and logistical consequences.

Provisioning each server in your data center with its own physical PCIe based SSD adapter would ordinarily be cost prohibitive. However, imagine pooling two or even four of these same resources in a centralized appliance, and then sharing them as a virtualized PCIe-based SSD to multiple server hosts! Now any server can have its own portion of an "in-demand" resource, and depending on the configuration, all server hosts can see the resource as a shared pool, though connected via PCIe.

With the shared resource model, the performance benefit alone is a huge improvement; for example, block copies within a shared logical device or data store, to support operations within a clustered database or virtual machine live migration between server hosts. The operation can be achieved at local speeds/feeds as if the storage resource was directly attached and local to the server host itself.

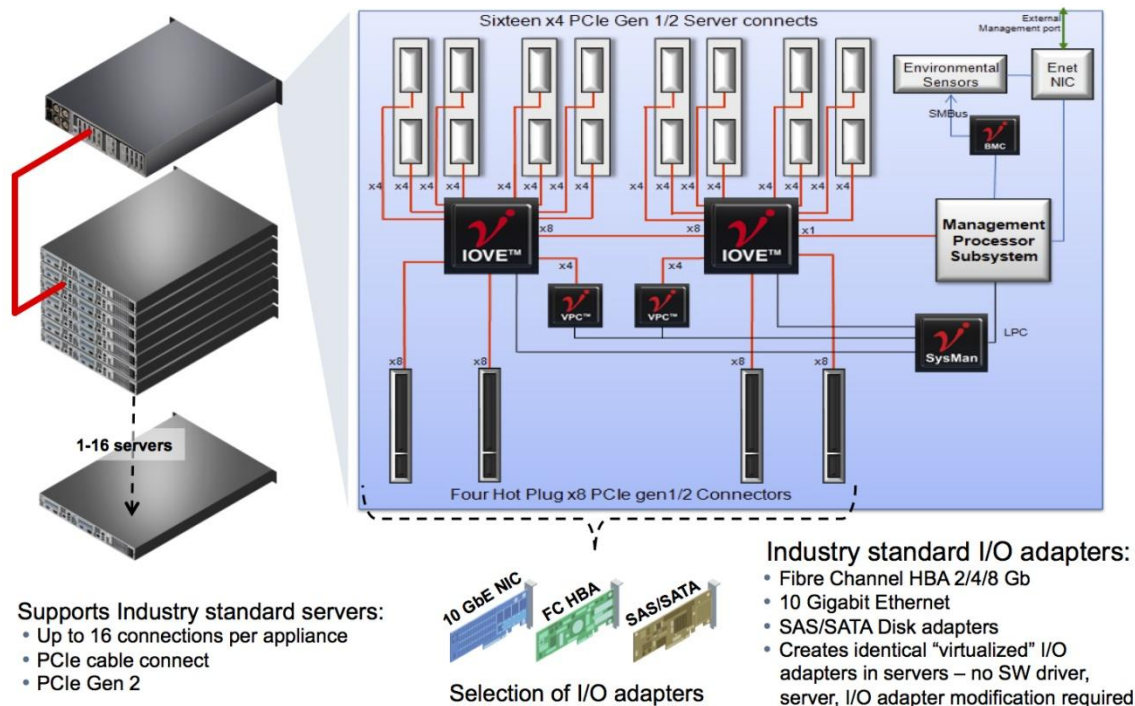
Utilizing PCIe sharing – smaller physical servers together with more densely populated racks gain the full advantages of fewer physical cables to manage, and fewer adapters in each server, but with all the right capabilities, in the amounts/types needed to provide appropriate performance/scale.

PCIe “Slot” constrained servers benefit from technology like Virtensys’ PCIe Sharing technology in the following ways:

- **Physical Layer Consolidation** PCIe sharing addresses the need to remove the physical and logical layer complexity, providing management agility.
- **I/O Density** Providing server hosts with greater amounts of bandwidth and differing types of I/O resources for densely populated server hypervisor hosts.
- **I/O Diversity** Servers become “agnostic” to the storage protocol used, as they can support all native environments and can easily migrate between them.
- **Hardware Acceleration** Server-to-server communication remains within the “zone” and I/O communications happen at the lowest latency and highest performance/bandwidth possible.
- **Resource Extension** I/O resources, such as traditional SSDs and PCIe SSDs can be leveraged through virtualization as an extension of memory capability; e.g. “host cache”.

Servers not only need access to the right type and number of I/O resources but also additional resource can come at a cost from an overhead perspective. The last thing an administrator needs is a server fully loaded with high performance interfaces that can’t keep up due to CPU/memory utilization.

The cost of processing traffic for a high performance network or SAN may mean sacrificing needed CPU cycles to be shared across virtual machines. With Virtensys’ PCIe Sharing technology, much of the PCIe “control” related traffic is offloaded from the server. This means that traffic related to processing communications at the PCIe layer is hardware accelerated within the Virtensys architecture.



Virtensys' architecture consists of two key components: the **IOVE** (I/O Virtualization Engine) and the **VPC** (Virtual Proxy Controller):

- The **IOVE** provide a 16 port PCIe multi-root aware switch; a high-speed switching fabric that provides 64 lanes of up to 320Gb/s of non-blocking bandwidth, fully compliant to the PCI-SIG MR-IOV specification.
The IOVE handles the separation of the control and data planes. This enables each I/O device to access multiple host memory spaces, effectively separating direct memory access (DMA) to data buffers and control structures.
- The **VPC** operates on the control path to provide the virtualization of multiple physical devices, and is a proxy between the I/O device and the connected server. The VPC can pre-fetch data from host memory in anticipation of a read operation from a device, or from device registers anticipating a read operation from a host. Since the VPC is implemented as a hardware state machine rather than through software and context switching, register access latency on a virtualized I/O device is comparable a physical I/O device, and in some cases can even improve the throughput of a virtual I/O device, compared to a physical device.

When the virtualized version of a physical device is presented to a server, it's as if the device were physically installed in the server. For example, a virtualized 10GbE adapter will be seen as a PCIe device in a server's BIOS, and a virtualized Fibre Channel HBA will go through its normal initialization processes during server boot – just like a physically installed adapter. In addition, these virtualized adapters require no additional Virtensys-proprietary drivers, nor is it necessary to boot into the hosts operating system or hypervisor to be able to operate the adapters. This makes features such as PXE (preboot execution) or boot from SAN a possibility.

Summary

Sharing in the data center has been proven technology approach for decades, seen as a better way to achieve scale, resource utilization and even improve performance. As we progress beyond what can already be shared (server, network, and storage) to other areas of the data center, we recognize that the time has arrived to leverage sharing for PCIe based resources.

Utilizing Virtensys' PCIe sharing technology – administrators no longer need to choose between physical server form factor. They can leverage smaller server footprints *and* gain the benefits of hyper-consolidation, *without* sacrificing I/O performance and scale. Furthermore they gain the benefits of I/O density and I/O diversity, with the added advantage of hardware acceleration and improved physical resource management.

For more information on Virtensys and how PCIe Sharing can work for you, check out: www.virtensys.com, www.facebook.com/virtensys or follow Virtensys on Twitter: [@virtensys](https://twitter.com/virtensys)