

Innovative I/O Virtualization  
Solutions for the Data Center



WHITEPAPER

## Accelerating 10 GbE Deployment with I/O Virtualization



Just a few short years ago, a typical server had a single high-performance CPU and one or two on-board 1 Gigabit Ethernet (1 GbE) NICs. How quickly things change...

The equivalent typical 1U or 2U server today will have two CPU sockets; each CPU containing up to six cores, and RAM capacity easily extending into hundreds of Gigabytes. This means that a single server today has more compute power than dozens of state-of-the-art servers as recently as five years ago.

What hasn't changed however is the network I/O. Despite the massive scaling in CPU and memory capacity, the same two 1 GbE NICs remain, meaning that the default network I/O connectivity bandwidth is the same as in those early servers. Furthermore, whilst the single CPU servers of yesterday were usually deployed for single or very few applications, today's high-performance, multi-core servers are called upon to serve several tasks – either by leveraging virtualization or simply sharing the performance and capacity in the OS (operating system).

What this means is that the performance bottleneck has not gone away – it has simply shifted away from the CPU – we now have an I/O bottleneck!

## 10 GbE to the rescue

If it was reasonable for yesterday's single CPU server to have two 1 GbE NICs, it follows that today's multi-core based servers require the full network bandwidth capacity that 10 Gigabit Ethernet (10 GbE) provides. For a number of years virtualization vendors have been recommending installing one or more quad-port 1 GbE NICs in order for the server's I/O capacity to keep pace with the CPU and RAM. However, while four or eight more Ethernet ports and cables might improve the I/O problem for an individual server, it creates problems elsewhere:

- CPU performance can be adversely affected by the trunking of multiple 1 GbE connections.
- The network infrastructure requires multiple Ethernet ports per server, together with the associated management of these ports.

It's for this reason that 10 GbE has been gaining interest for new server deployments. However, deploying a 10 GbE network infrastructure is not a trivial task, and the cost of deployment in a data center can be considerable. Some critical decisions confront the server and network administrators when evaluating a new 10 GbE networking infrastructure:

- **Cabling:** While it is possible to use familiar RJ-45 connectors and copper for 10 GbE, this is far from ideal for many implementations. Alternative short-range, low-latency twin axial copper cabling and short and long range optical fiber are more appropriate alternatives for many circumstances.

- **Oversubscription rates and aggregation ratios impacting network topology:** If a server requires 10 GbE of network I/O, even if only occasionally, how does the IT infrastructure adapt and deal with the requirements of multiple servers requiring 10 Gigabits of network bandwidth?
- **Ensuring proper design localizes much of the network traffic:** minimizing the impact on the rest of the network, and reducing the cost and complexity of deploying 10 GbE throughout the data center.

## Using I/O Virtualization to manage I/O provisioning

I/O Virtualization (IOV) works by allowing I/O adapters to be separated, consolidated and virtualized away from the physical confines of a server enclosure. These are then shared back to the same physical servers in the form of virtual I/O resources. Utilizing a solution such as Virtensys' VIO-4000 Series IOV Appliances, server administrators are able to provision, manage and optimize I/O resources across entire groups of servers – as opposed a server at a time.

Dependent on model, the VIO-4000 IOV Appliance consolidates and virtualizes 10 GbE NIC, RAID storage, and Fibre Channel Host Bus Adapter (FC HBA) connectivity. It connects to the native PCIe bus in the server using a single industry-standard PCIe (PCI Express) cable per server. The appliance provides fully switched non blocking PCIe sharing of the physical NICs and FC HBAs to up to 16 servers, in the form of virtual I/O which behaves exactly as if physically installed server I/O adapters, requiring no changes to the servers, OS, applications or I/O adapter device drivers.

Developed specifically for 10 GbE deployments, the VIO-4004 IOV Appliance is installed with eight 10 GbE uplink ports providing a total of 80 Gigabits of bandwidth that can be centrally managed, optimized, and shared amongst the attached servers.



**Figure 1 – VIO-4004 10 Gbps Ethernet I/O Virtualization Appliance**

Virtensys' VIO-4000 IOV Appliance connects to industry-standard x86 servers using standard PCIe cables. Each connection provides 20 GT/s (GigaTransfer per second). This equates to 16 Gbps of pure PCIe bandwidth. In addition, multiple server to appliance connections can be added per connected server providing high performance and/or high availability.

## The VIO-4000 Series in 10 GbE Deployment

Virtensys' VIO-4000 Series IOV Appliances allow server administrators to configure their servers as generic high-performance computing resources, and connect to them using their native PCIe bus – allowing them to “wire them once – and provision at will”.

The IOV Appliance becomes an extension of the PCIe bus of the connected servers, and the 10 GbE NICs in the appliance (up to 8 ports of 10 GbE per IOV appliance) – as well as other I/O resource types (FC HBAs and Virtual Direct-Attached RAID storage) – can be configured, provisioned and managed among all the connected servers using the browser-based Virtensys Management Suite (VMS).

For example: using the VIO-4004, a server administrator can connect up to 16 server connections once and dynamically re-configure the servers' I/O via the VMS without making any physical changes to the servers. The 80 Gigabits of network I/O availability can be dynamically allocated and shared between connected servers.


A guaranteed minimum network bandwidth can be assigned per virtualized 10 GbE interface for each server connection. If more bandwidth is available and the server requires it, the appliance dynamically allocates the incremental bandwidth where it can be used.

## Optimizing Server I/O

Using the VIO-4000 IOV Appliance it is possible to configure high-performance, multi core servers with **no** local internal I/O. The 10 GbE NICs are configured, optimized, managed and shared from the appliance – transparently to the applications, OS or server virtualization hypervisors. I/O provisioning can be managed across groups of servers rather than a single server at a time.

## Problems with traditional approaches

When configuring a high-powered server today, it is common to find the native network capacity of the server has become a bottleneck – particularly where server-consolidation initiatives increase the utilization of the CPU capacity.



**Figure 2 – A typical 2U, 2 Socket, 12+ Core Server Configured for Virtualization**

As the typical example in Figure 1 shows, the traditional approach to reducing the network bottleneck has been to add more and more 1 GbE NICs. Whereas a single 1 GbE NIC was appropriate for a typical server five years ago, today it is common to find 4 or 8 GbE ports or even more required for a single server, owing to the increase in compute power and capacity. Adding NICs to each server simply moves the problems elsewhere:

- Eight or ten access/Top of Rack) network ports per server are required.
- The commensurate power, reliability and cooling problems associated with multiple NICs and cabling.


A single traditional 10 GbE NIC in each server could mitigate these problems but, in turn, raises its own:

- Each server would require its own 10 GbE port, at considerable expense in switching gear in addition to the NICs and cabling. If the data center is not fully 10 GbE ready it may mean negotiating with network administrators for scarce 10GbE ports.
- Adding bandwidth to a server must be done in 10 GbE increments – a very coarse level of granularity.

### Leveraging scarce 10GbE resources

As can be seen in Figure 3, utilizing the Virtensys VIO-4004 IOV appliance, connected to up to eight 10 GbE ports in the data center, network resources to all the connected servers are optimized and managed centrally. It is also possible for additional servers to be added without having to request additional network ports.

The VIO-4004 allocates a guaranteed minimum bandwidth to each server, and yet the servers can still dynamically access additional network capacity if available.



**Figure 3 – An Equivalent Performance 1U Server with a Virtensys PCIe Server Connect Adapter**

### Enterprise-class capabilities

If maximum Ethernet bandwidth is required, up to four dual-port 10 GbE NICs (a total of 80 Gigabits of throughput) are available to be shared amongst up to 16 servers from a single VIO-4000 IOV Appliance.

Virtensys VIO-4000 Series IOV Appliances come with hot-swappable fans and can include hot-swappable redundant power supplies. If even higher availability is required, multiple VIO-4000 Series IOV Appliances can be deployed for a fully-redundant implementation. In this configuration, a total of 160 Gigabits bandwidth can be provisioned to up to 16 connected servers while both VIOs are active.

### Leveraging I/O Virtualization in non-virtualized server environments

The elegance of an I/O Virtualization solution such as Virtensys' is a logical extension to the benefits of server virtualization (VMware<sup>®</sup>, Hyper-V). Yet I/O Virtualization also even provides benefits for those hard-to-virtualize applications.

For example: High Performance Computing (HPC) applications are seldom virtualized, owing to the need for maximum performance and reduced networking latency. Using the VIO-4000 IOV Appliance, a system administrator can leverage the I/O hardware virtualization of the Appliance to allow each of up to 16 servers attached to a single VIO-4000 IOV Appliance to 'see' four dual-port 10 GbE NICs and up to 32 hardware-virtualized NICs – regardless as to whether they are using virtualization.

For clusters of compute nodes, the appliance provide higher performance and lower overhead than can be realized with add-in NICs and external switching.

### Transparency – within the server and to the network

All of this power and flexibility comes transparently: The server and the OS or hypervisors see the vNICs (and/or HBAs, RAID controllers) as if they were locally installed in the server rather than virtualized in the IOV appliance. The server simply discovers the new interface, loads the native device driver and starts to take full advantages of its new I/O adapter resources.

Innovative I/O Virtualization  
Solutions for the Data Center

**About Virtensys**

Virtensys develops industry-leading, patented PCI Express®-based I/O Virtualization technologies for servers and storage platforms, revolutionizing the way I/O infrastructures are deployed and used in data centers, and delivering significant improvements in I/O utilization, cost, performance, power consumption, and management.

The deployment of the IOV switches is totally non-disruptive and dramatically reduces IT complexity and expenses.

[www.virtensys.com](http://www.virtensys.com)



**US Headquarters**

Virtensys Inc.  
14908 NW Greenbrier Parkway  
Beaverton  
OR 97006  
USA  
Tel: +1 503-210-5190  
Fax: +1 503-533-5707  
Email: [info\\_request@virtensys.com](mailto:info_request@virtensys.com)

**European Headquarters**

Virtensys Ltd  
5500 Lakeside  
Cheadle  
Cheshire  
SK8 3GR  
United Kingdom  
Tel: +44 (0)161 495 1530  
Fax: +44 (0)161 491 4897