



Industry Brief
**Renaissance in VM
Network Connectivity**

Featuring

An approach to network design
that starts with the server

Executive Summary

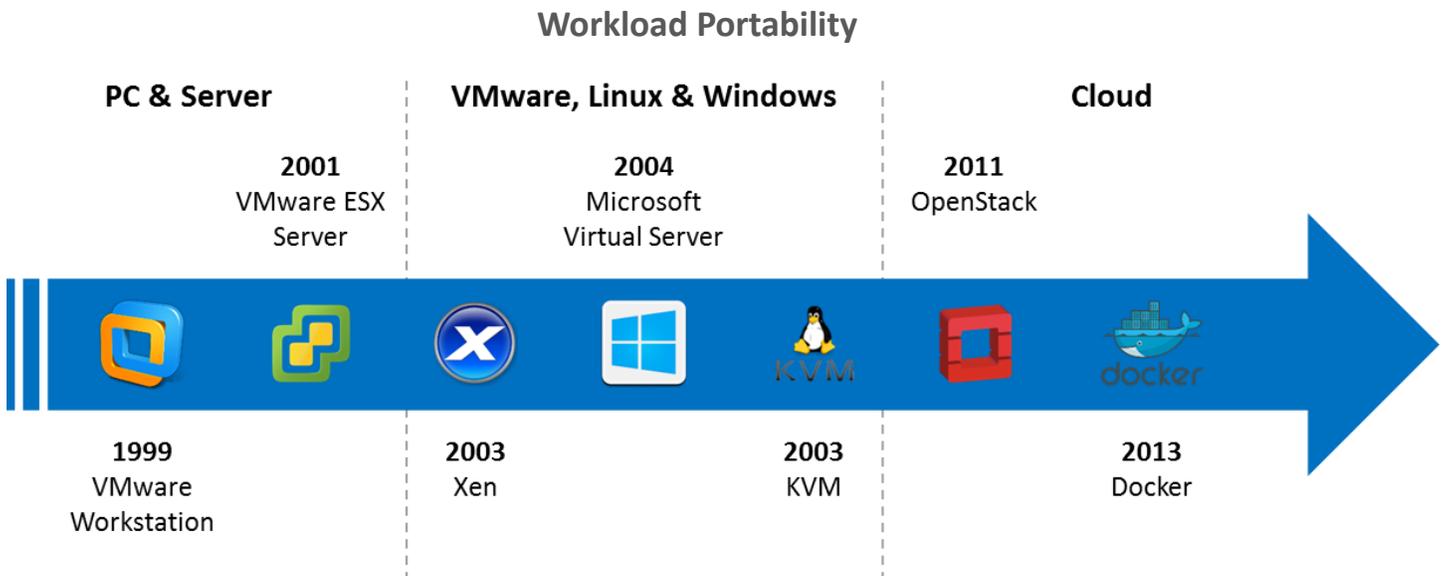
Server Virtualization is Leading the Way to Software Defined Data Centers

This industry brief is about a renaissance in how virtual machines are networked. To help the reader understand why the industry is moving out of the dark ages, we first examine the driving force behind the new technology: hyperscale service providers and software define data centers.

Necessity is the mother of invention, and cloud service providers invented dozens of new information technologies to operate profitably while deploying millions of virtual machines dispersed geographically around the world. Virtual networking for VMs is pervasive in hyperscale data centers because it is a core capability for meeting requirements for cost-effective server virtualization at scale.

Renaissance in VM Network Connectivity in the Enterprise

The same virtual networking technology Amazon and Microsoft Azure use to achieve QoS, greater workload portability and lower cost, is now available to the masses. Read this report to learn more about how to design more efficient network connectivity for your virtualized servers.

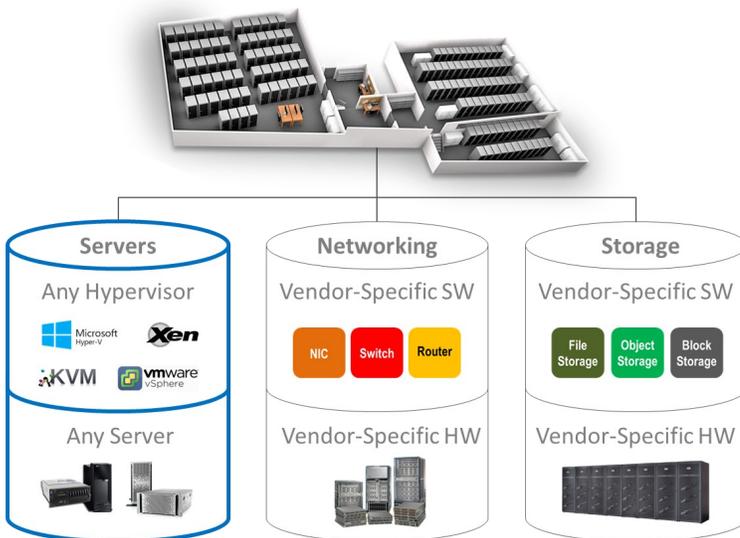


Originally invented to increase the utilization of server hardware in a single data center, server virtualization now underpins expanding workload portability for computing across geographically dispersed cloud data centers.

The Model for Data Center Efficiency

Software Defined Data Center: Server Virtualization Leading the Way

Approximately 80% of new workloads deployed in enterprise data centers are hosted on virtual machines, otherwise known as software defined servers. Hyperscale data centers have taken that concept two steps further with software defined storage and software defined networks. Almost every enterprise IT organization is at some stage of transforming their infrastructure into private clouds with a software defined data center architecture allowing them to efficiently manage virtual resources across multiple data centers.

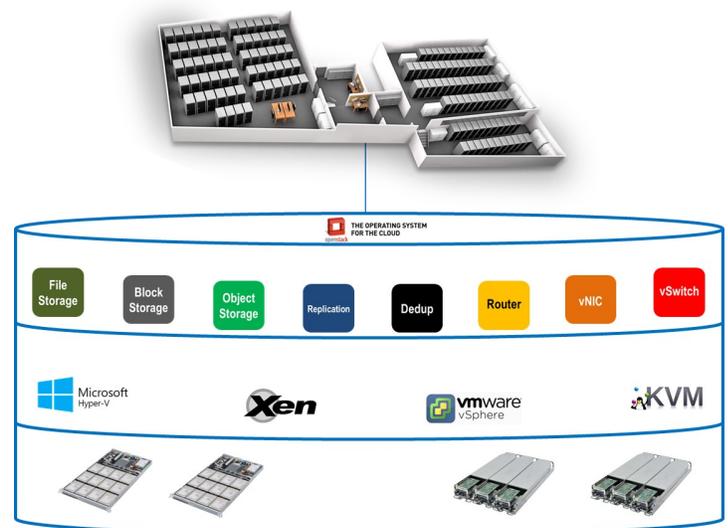


Traditional Data Center

In a typical enterprise data center, IT organizations have broken free from server vendor lock-in with hypervisors which can virtualize any server using an x86 processor. However, most networking and storage environments remain silos of vendor-specific hardware and software.

Software Defined Data Center

In software defined data centers, IT organizations have transformed their networking and storage infrastructure from vendor-specific hardware into open-source based software and hardware. This is accomplished by deploying apps on virtual machines and industry standard servers—all under the centralized control of a cloud operating system such as OpenStack.



Today, over 80% of new workloads in the enterprise are deployed on virtual machines, and the benefits of server virtualization have signaled IT organizations it's time to consider a migration to software defined networking and storage.

Trending Now in the SDDC

10GbE Server Connectivity

The penetration of server virtualization has reached the point where almost all new workloads are running on virtual machines, IT pros are comfortable with tier-1 apps running on virtual machines, and VM density is increasing to the point where network I/O has become a critical issue.

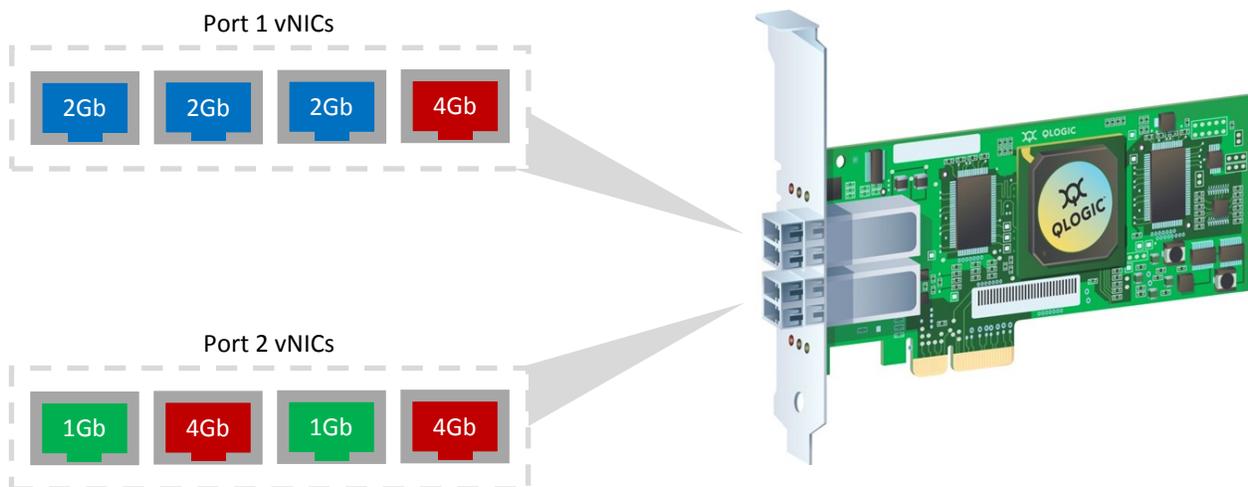
This aggregation of VMs, workloads and I/O explains why, after ten years on the market, the deployment of 10GbE network connectivity has become a necessity. This trend is also driving adoption of virtual networking technology to partition 10GbE ports and shape the virtual network links on a per-VM basis.

Past Virtualized Server Network Connectivity—6 x 1GbE Ports



According to IT Brand Pulse surveys, a typical virtualized server has an average of six 1GbE ports configured and the only network virtualization deployed is teaming. IT pros say the biggest issue with 1GbE is lack of performance, and the biggest issue with upgrading to 10GbE is cost.

Emerging Virtualized Server Network Connectivity—10GbE NICs Partitioned into vNICs



What's trending in server network connectivity is the use of 10GbE NICs in combination with partitioning. The two together can provide more bandwidth than a pair of teamed NICs, and offset the cost of 10GbE by replacing up to ten 1GbE NICs.

Adoption of Virtual 10GbE Networks

Less than Half of IT Pros are Virtualizing Network Connectivity to VMs

While the use of virtual NICs and virtual switches is pervasive in hyperscale data centers, less than half of enterprise IT pros surveyed said they will deploy these virtual networking technologies in 2015.

This indicates there is an opportunity for many enterprise IT organizations to lower the cost of their virtualized servers by partitioning physical networks into many virtual networks, and to improve the performance of their applications by implementing QoS to guarantee bandwidth to select virtual networks and workloads.

More than Half Prefer to Use a Hypervisor for Virtual Network Management

Over half of IT pros surveyed indicated they would prefer to use a hypervisor to manage their virtual networks, almost twice as many who preferred to manage virtual networks through a switch.

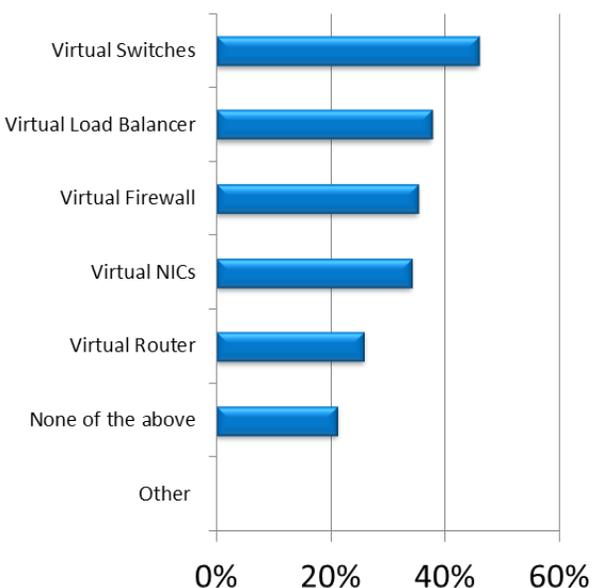
This indicates a preference for a virtual networking management platform which is common to all servers.

All are Balancing the Benefits of Virtual Networks with Available Resources

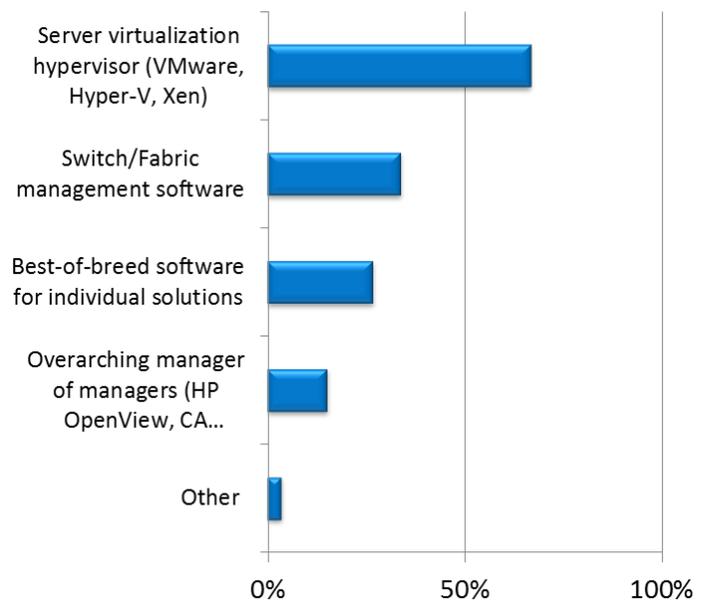
Implementing virtual networks is not trivial. Busy IT pros continue to map physical network ports to their VMs because they don't have the time and/or budget to learn the technology and evaluate products. Even large enterprises with thousands of VMs must balance the benefits of virtual networks with available resources.

New approaches to virtual networking are needed which will make adoption easier for small enterprises, and lower the cost for large enterprises.

Which of the following network virtualization products do you expect to evaluate or deploy in 2015?



I want to manage my virtual network resources from :

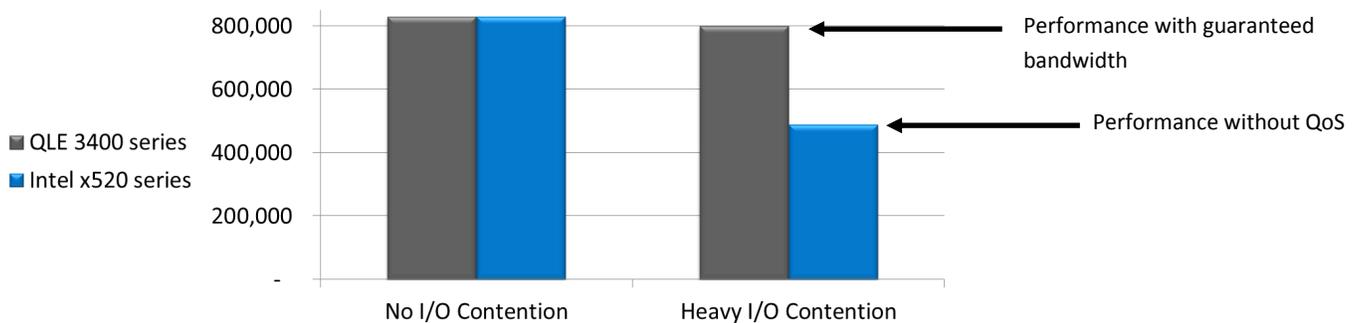


Killer Apps for Virtual Networks

Guaranteed Quality of Service for Specific VMs and Workloads

Dedicating physical 10Gb network ports to VMs can get expensive. That's why a killer app for virtual networks is to partition the physical ports and guarantee bandwidth to the most important workloads. The chart below shows the performance of a SQL workload, which was sharing a physical 10GbE LOM port, was degraded when there was heavy I/O contention. In contrast, performance was maintained when a QL3400 NIC was configured to prioritize SQL traffic on one partition, guaranteeing bandwidth to the SQL workload.

SQL Transactions per Minute with & without Network Partitioning (NPAR) QoS



Fine Tuning Virtual Networks by Traffic Type

Shaping virtual networks by making use of granular controls for ports and resource pools can further increase application performance and availability. The diagram below shows examples of how server admins can configure virtual ports for predefined resource pools with VMware ESX Network I/O Control (NIOC)

VMware ESX Network I/O Control

Admins may configure virtual port groups with these capabilities:

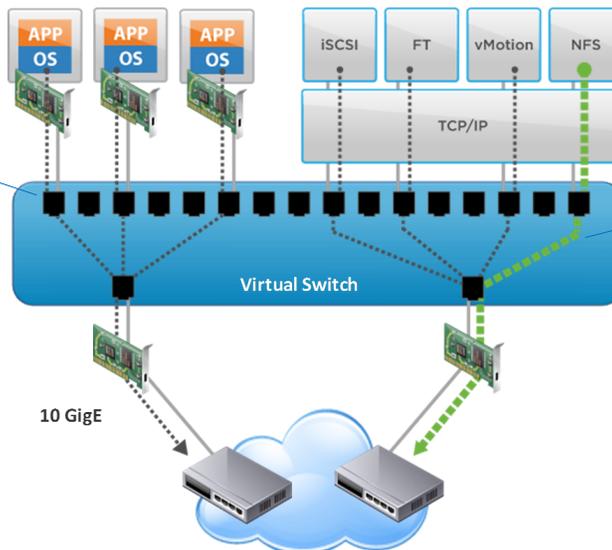
Isolation - Ensures traffic isolation so that a given flow will never dominate over others, preventing drops.

Shares - Allows flexible networking capacity partitioning to help users deal when flows compete for the same resources.

Limits - Enforce traffic bandwidth limit on the overall VDS set of dvUplinks.

Load-based Teaming - Efficiently uses a VDS set of dvUplinks for networking capacity.

IEEE 802.1p tagging - Tags outbound packets from the vSphere host for proper handling by physical network resources.



NIOC classifies traffic into these predefined resource pools:

- VMware vSphere vMotion
- iSCSI
- FT logging
- Management
- Network File System (NFS)
- Virtual machine traffic
- vSphere Replication traffic
- User defined

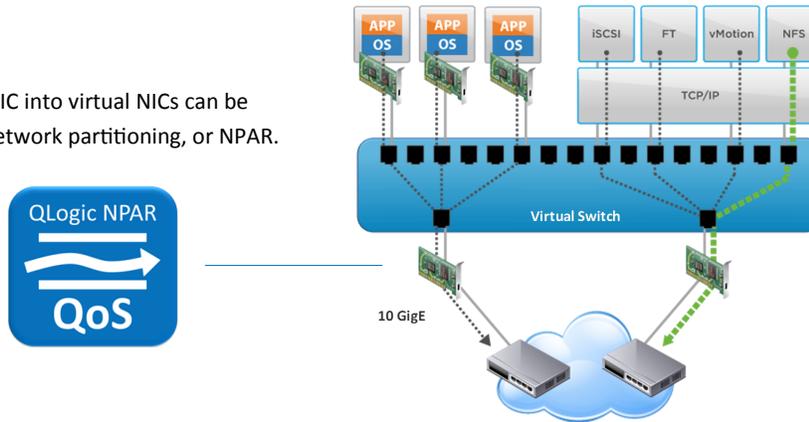
4 Virtual Networking Solutions

1

Use *any* Network Partitioning Tool to Consolidate NIC and Switch Ports, and to Guarantee Bandwidth to Critical Apps (QoS)

Considering that about two thirds of enterprises have not implemented vSwitches and vNICs, most enterprises can lower the cost of their virtual server network connectivity and improve application performance by taking the first step with their choice of switch, hypervisor, operating system or NIC tools for basic partitioning and QoS.

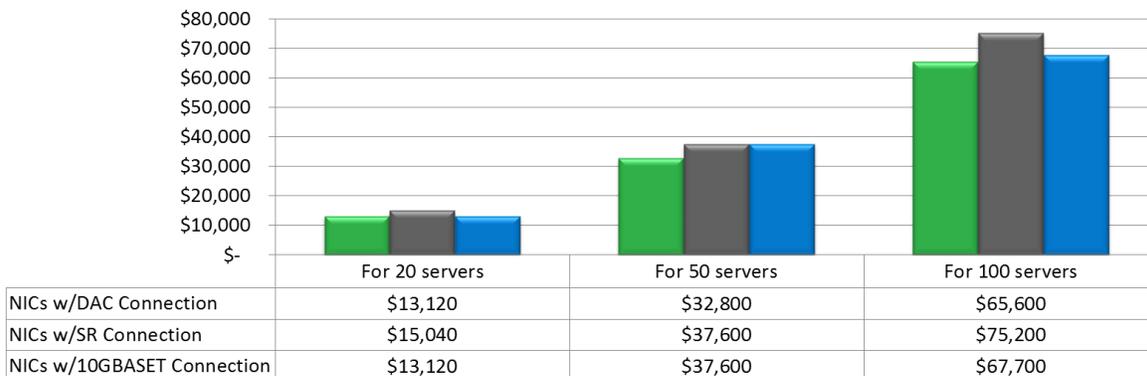
Carving a physical NIC into virtual NICs can be done with QLogic network partitioning, or NPAR.



Cut Costs by Sharing High-Performance NIC and Switch Ports

It is common for server racks with high VM densities to use high-performance 10GbE NIC and switches to ensure there is sufficient bandwidth. The quantity and cost of 10GbE links can be reduced by virtualizing ports to increase utilization. The chart below shows that for 50 servers connected via 10G Base-T, a savings of \$75,200 can be realized by increasing utilization of NICs resulting in elimination a single NIC per server.

Cost Reduction by Eliminating One 10GbE NIC per Server

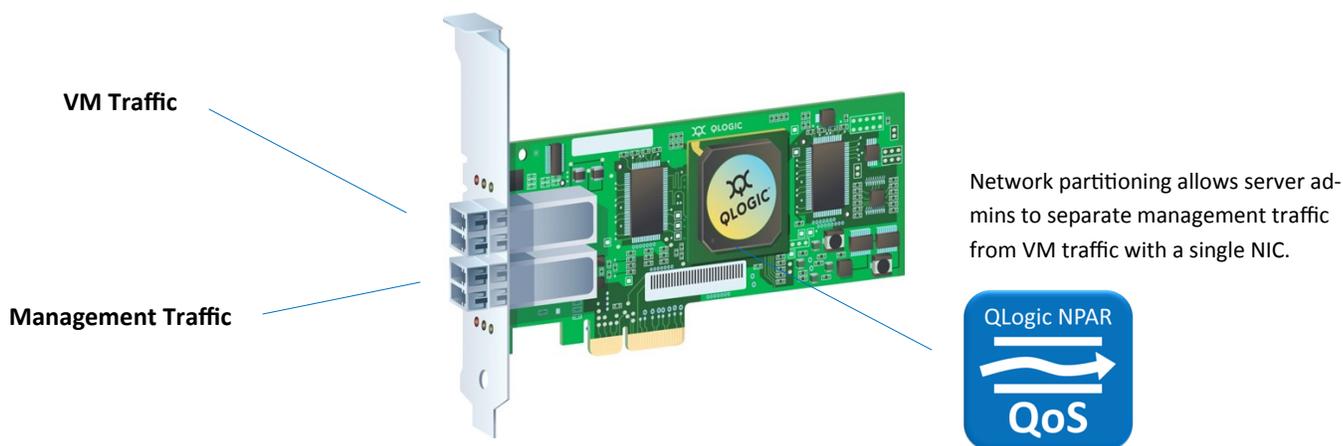


4 Virtual Networking Solutions

2

Keep all NIC Ports Available for Data Traffic with Hyper-V

There are numerous [articles](#) about Hyper-V networking best practices. One best practice which is mentioned frequently is that you should separate Hyper-V QoS management traffic from Virtual Machine traffic on the physical network.

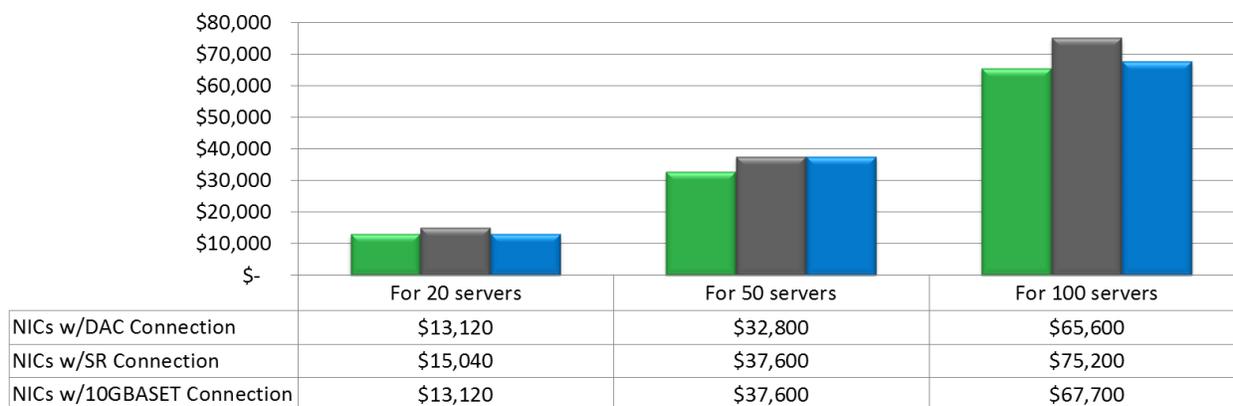


Network partitioning allows server admins to separate management traffic from VM traffic with a single NIC.

NIC-based QoS Enables Partitions on All Ports

A new best practice is to use NIC-based network partitioning (NPAR) for QoS. Because NIC firmware is not accessible in the same way as operating system resident software, you can mix VM and QoS traffic without compromising network security. You can also buy-back high-performance network ports which otherwise would be dedicated to management traffic. You could achieve a cost savings of over \$30,000 by implementing this new best practice.

The Cost of Giving Up 1 Port (Adding one 10GbE NIC)



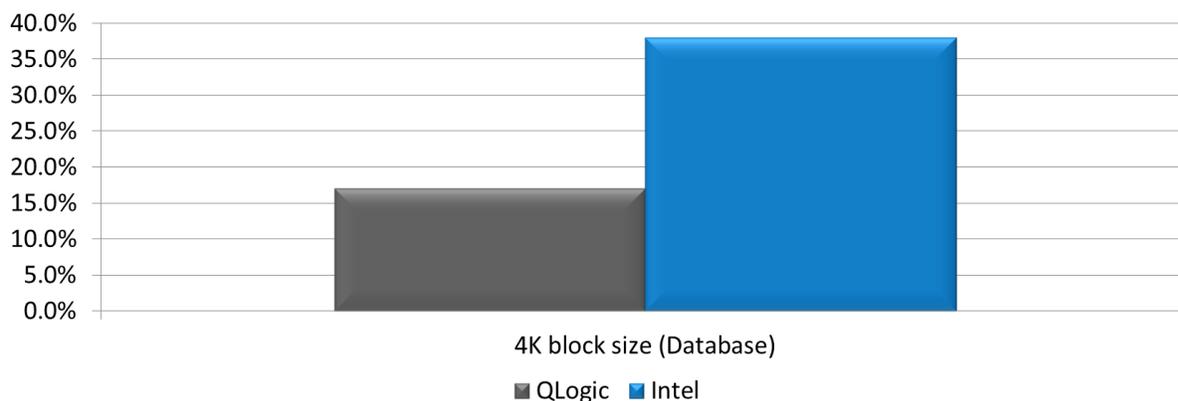
4 Virtual Networking Solutions

3

Keep CPU Resources Available for VM Workloads

Many server administrators underestimate the amount of CPU require to process network protocols which are not offloaded by some NICs. For example, NICs from Intel which use server CPU to “on-load” iSCSI, VXLAN and NVGRE protocol processing, can steal a third of your server’s processing power.

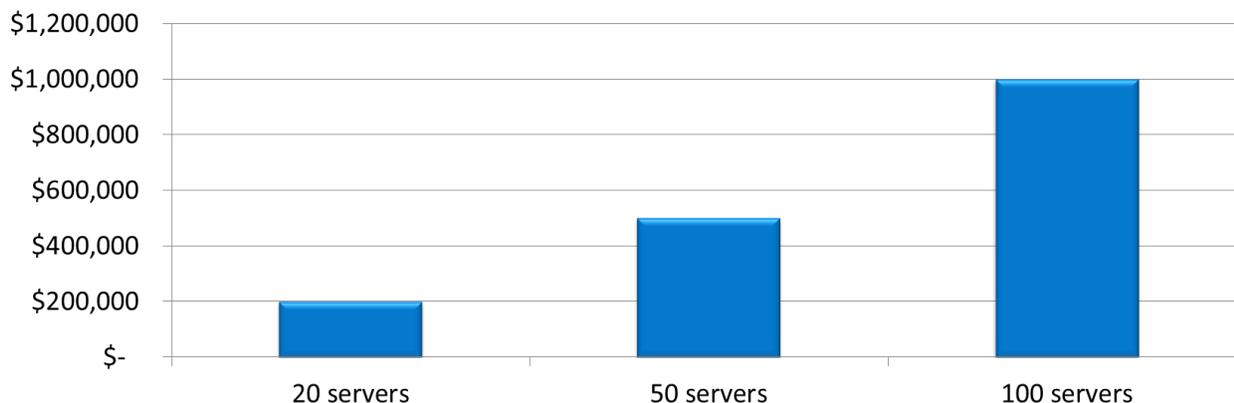
CPU Utilization with iSCSI Off-load (QLogic) vs. On-load (Intel)



The Cost of On-Loading Network Management & Protocol Processing

The more expensive your server, the more expensive it is to on-load network protocol processing. The chart below shows the cost of on-loading the iSCSI processing shown above for 20, 50 and 100 servers which each cost \$20,000.

The Cost of On-Loading Network Protocol Processing with \$20,000 Servers



4 Virtual Networking Solutions

4

Consolidate Virtual Networking

It is estimated that only 15% of IT organizations who are using hypervisors, use a single hypervisor. In addition, most IT organizations which have deployed virtual networking, are using virtual networking tools from their switch vendors. One way to improve the efficiency and lower the total cost of your virtual networks, is to consolidate your virtual networking tools where possible.

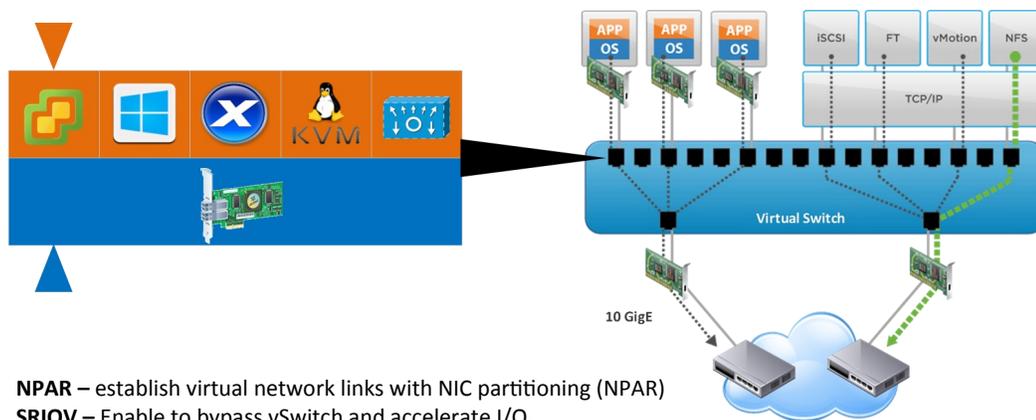


By Layering Virtual Networking Services on a Common Platform

A new approach for connecting VMs with virtual networks is to layer services on NIC-based partitions. This approach allows admins to standardize on a set of virtual network partitioning tools which are hypervisor and switch independent, and whose virtual network settings migrate seamlessly along with the VMs.

Software-based virtual networking services

- Isolation** - Ensures traffic isolation so that a given flow will never dominate over others, preventing drops.
- Shares** - Allows flexible networking capacity partitioning to help users deal when flows compete for the same resources.
- Limits** - Enforce traffic bandwidth limit on the overall VDS set of dvUplinks.
- Load-based Teaming** - Efficiently uses a VDS set of dvUplinks for networking capacity.
- IEEE 802.1p tagging** - Tags outbound packets from the vSphere host for proper handling by physical network resources.



Hardware-based virtual networking services

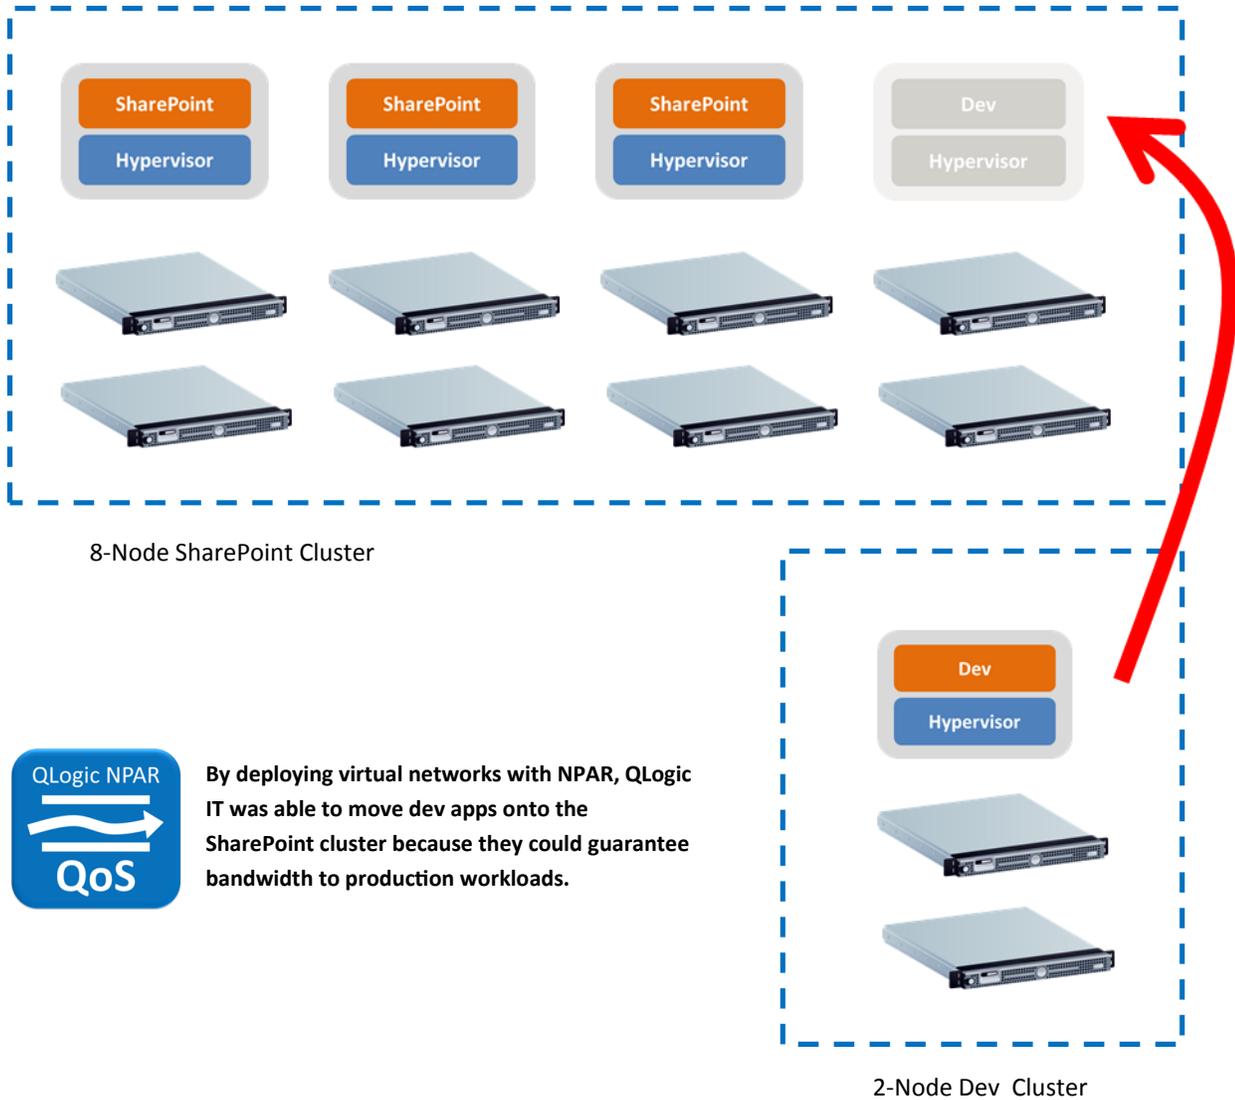
- NPAR** – establish virtual network links with NIC partitioning (NPAR)
- SRIOV** – Enable to bypass vSwitch and accelerate I/O.
- HW Offload** – Network protocol processing is offloaded from the server for partitioning, SRIOV, iSCSI storage, and VXLAN tunneling
- Hypervisor and Switch Independent** – Standardize on one network virtualization scheme for a data center

Real Life Case Study

QLogic Eats its Own Dog Food

Until recently, QLogic IT knew they had extra processing power in their 8-node SharePoint cluster, but did not have a way of using the processing power without risk of degrading SharePoint performance. The IT solved the problem by deploying the company's own NIC-based network partitioning (NPAR). NPAR virtual network allows the CPU resources to be shared by production and dev applications, while guaranteeing bandwidth to the production apps. The result is QLogic IT has one less cluster and two less servers to manage and keep under service contracts.

QoS Allows Dev Apps to Migrate to SharePoint Cluster



QLogic FastLinQ Ethernet Adapters

#2 in 10GbE NICs for the Enterprise

QLogic has a broad portfolio of high-performance Ethernet adapters for enterprise servers—all of which are designed to deliver a rich set of virtual networking services for virtual machines. For basic LAN and NAS connectivity, the company offers the 3200 Series and 3400 Series. If you want to consolidate LAN and SAN connectivity to servers, QLogic offers the 8300 Series and 8400 Series converged network adapters with full offload of the CPU-hungry iSCSI, FCoE or Fibre Channel protocols. The 8300 Series can even be converted to a native 16Gb Fibre Channel HBA.

#1 at 25GbE

QLogic has publically stated that OEMs are qualifying their 25G and 100G converged network adapters, so I fully expect QLogic to be first to market in 2016 with products which are generally available from OEMs. The new adapters will open up server connectivity for applications that need more bandwidth, and will be a cost-effective alternative to multiple 10G adapters as the price of 25 G optical transceivers and cables drop.



3200 Series	3400 Series	3500 Series	3600 Series	8300 Series	8400 Series
					
10Gb NIC	10Gb NIC	25G CNA	100Gb CNA	10Gb CNA	10Gb CNA
Available Today	Available Today	Coming in 2016	Coming in 2016	Available Today	Available Today
High-performance 10G NIC.	Available in 10G BASE-T versions and features hardware offloads for NVGRE and VXLAN tunneling protocols used in private clouds.	World's 1st 25Gb CNA. Serves as a NIC in 25Gb LANs and as a HBA in iSCSI and FCoE SANs.	World's 1st 100Gb CNA. Serves as a NIC in 100Gb LANs and as an HBA in iSCSI, and FCoE SANs.	For LAN and iSCSI, FCoE SANs. Serves as a NIC in LANs and as an HBA in iSCSI, and FCoE SANs. Can be converted to a 16Gb Fibre Channel HBA.	For LAN and iSCSI, FCoE SANs. Serves as a NIC in LANs and as an HBA in iSCSI, and FCoE SANs.
					



VM Networking Checklist

For typical IT organizations using hypervisors from VMware and Microsoft, the following is a checklist of things to consider before implementing or optimizing virtual networks to your virtual machines. Each item represents an opportunity for you to deploy server connectivity more efficiently and cost effectively.

(If You Haven't Already) Implement Virtual Networks

The biggest opportunity to impact efficiency and reduce the cost your server networking is implement *some* type of virtual network partitioning tool in the NIC, or hypervisor vSwitch. If sharing 10GbE ports can save one NIC and switch port per server, the savings will add-up fast.

When 25GbE, 50GbE and 100GbE NIC ports emerge in 2016, implementing network partitioning will be required to making full use of the bandwidth available on each port.

Use NIC-based Partitioning as a Virtual Networking Platform

If you really want to be efficient, consider using NIC-based partitioning because it is hypervisor and switch independent. It will allow you to standardize on one partitioning tool which works with any hypervisor and vSwitch.

Offload Network Protocol Processing

CPU vendors will tell you that on-loading network protocol processing is not a big deal, because they sell processors. It really is a big deal because it can silently eat-up a third of the server CPU you bought to support VMs and their workloads. Use NICs which offload iSCSI, FCoE, NVGRE and VXLAN and you will preserve your investment in processing power for business apps. This issue is especially important if you are using iSCSI or FCoE interface storage which is I/O-intensive. Don't giveaway a third of your CPU power if you can avoid it by simply deploying a NIC with offload.

Summary

The Bottom Line

The next wave of server virtualization for hyperscales centered around virtualizing the network connections to better scale and to lower costs. Those benefits are now available to the enterprise. Surprisingly, the humble NIC plays a key role in the new era of software defined data center because it is the gateway to the industry standard servers which host business and infrastructure apps, including NFV and software defined storage.

The bottom line is that selecting your NIC technology has suddenly become a strategic decision if you want to wring out the efficiencies and cost savings enjoyed by the big hyperscales. Don't settle for fast and dumb on-load NICs. Look for fast and smart NICs which are designed from the ground up for virtual networking.

Recommendations

- Estimate how much time and capital you would save if you implemented virtual networks.
- Create a virtual networking strategy that extends from inside server to other data centers.
- Standardize on a NIC platform which supports network partitioning (NPAR) and layer other services on top.
- Start small and fail fast. Deploy QoS in a DevOps lab or in a non-critical production environment.

Resources

www.2550100.com

[vSphere Distributed Switch options](#)

[Hyper-V networking best practices](#)

[The 2550100GbE Alliance](#)

[QLogic Intelligent Ethernet Adapters](#)

[QLogic Converged Network Adapters](#)

[IT Brand Pulse](#)

The Author



Frank Berry is founder and senior analyst for IT Brand Pulse, a trusted source of testing, research and analysis about cloud IT infrastructure, including operating platforms, servers, storage and networking. As former vice president of product marketing and corporate marketing for QLogic, and vice president of worldwide marketing for the automated tape library (ATL) division of Quantum, Mr. Berry has over 30 years experience in the development and marketing of IT infrastructure. If you have any questions or comments about this report, contact frank.berry@itbrandpulse.com.