Intelligent Queueing Technologies for Virtualization

A Dell-Intel-VMware Ecosystem Solution to Enhance Performance

Introduction

The maturation of virtualization technologies has brought with it an ecosystem of servers, network I/O silicon, and virtual machine (VM) hypervisor vendors whose products are engineered specifically to complement one another. Components developed under those collaborative relationships enable smooth interoperation for solution stacks that usher in a new era of virtualized efficiency and performance. Relationships among ecosystem participants can help to relieve the burden of solution integration from end customers, since many engineering challenges have been addressed by the hardware and software makers long before the solution components reach the market.

These benefits are especially important to get the most out of new technologies, since infrastructure-wide challenges often benefit from being addressed by multiple parts of the solution stack. One such challenge is the increasing amount and complexity of network I/O that is generated as the number of virtual machines (VMs) being supported per physical server increases. Intel® Virtualization Technology for Connectivity[1] (Intel® VT-c) is a suite of network I/O technologies that helps in reducing the network I/O bottleneck on the virtualized server. One concern that arises with high levels of virtualized traffic is the packet-sorting burden on the hypervisor, which can create a significant bottleneck to overall system performance.

Virtual Machine Device Queues[2] (VMDq) is one such technology under Intel VT-c that offloads network I/O data processing from the virtual switch in the hypervisor to the network I/O silicon. As a result, servers that implement VMDq can improve I/O throughput for faster, more efficient networking in virtualized environments. The complete solution benefits from the capabilities of components provided by multiple ecosystem participants:

- **Intel® Gigabit and 10 Gigabit Ethernet Server Adapters** that implement VMDq are available for both fiber and low-cost copper infrastructures.
- **VMware NetQueue** provides an interface in VMware ESX* for the customer to achieve the benefits of VMDq on its servers.
- **Dell PowerEdge* Servers** complete the ecosystem, enabling customers to enjoy the benefits of VMDq by running the VMware ESX hypervisor solution along with Intel 10 Gigabit Ethernet Server Adapters.

Increasing Virtual Machine Density Puts Bigger Burdens on Hypervisors

Deploying virtualized environments on more powerful platforms is a growing practice among IT departments in order to consolidate server workloads and reduce data center footprints. However, this practice can have a significant impact on system and application performance as workloads become increasingly bound by network I/O. While IT managers are adding greater processing power and reducing the infrastructure footprint, this kind of consolidation does not necessarily mean more efficient network throughput in the virtual environment. A balance between system performance and networking capabilities is required to achieve optimal application services from consolidation.

vmware® intel® DELL

In conventional virtualized environments, a virtual switch in the hypervisor manages network I/O activities, including packet sorting, to negotiate traffic exchange between multiple VMs and the physical server adapter. With more VMs and increased traffic through the platform, the hypervisor requires more CPU cycles to sort data packets and route them to the destined VM (Figure 1), reducing the amount of CPU capacity available for applications. Intel's VMDq is a breakthrough technology that helps to reduce the burden on the hypervisor while improving network I/O performance through the virtualized platform.

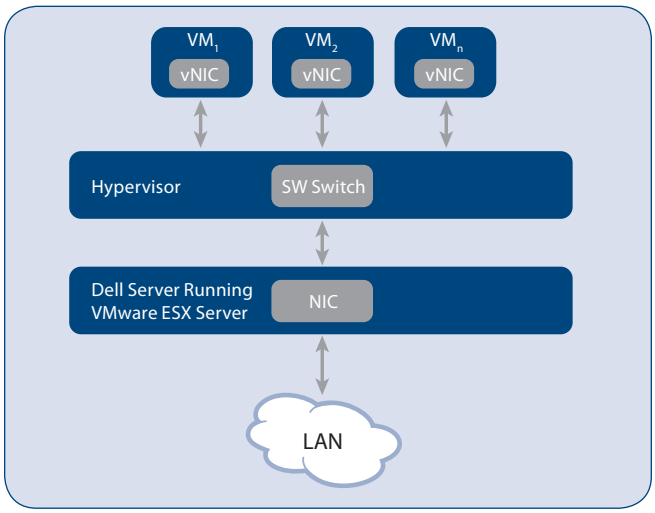

Figure 1. In conventional virtual environments, the hypervisor manages network I/O, creating a potential bottleneck as the number of packets to be sorted increases.

Advanced Queueing Technologies Reduce I/O Bottlenecks

IVMDq is a network I/O silicon technology that offloads the network I/O data packet sorting burden from the hypervisor. Multiple queues and sorting intelligence in the network I/O silicon help in directing the I/O traffic flow to the destined VMs more effectively and efficiently, as shown in Figure 2. This frees up processor cycles for application work and increases overall system performance.

Receive Packets

As data packets arrive at the network adapter, a Layer 2 classifier/sorter in the network controller sorts and determines which VM each packet is destined for based on MAC addresses and VLAN tags. It then places the packet in a receive queue assigned to that VM. The hypervisor's switch merely routes the packets to the respective VM instead of performing the heavy lifting work of sorting data. Thus, VMDq improves platform efficiency for handling receive-side network I/O and lowers CPU utilization for application processing.

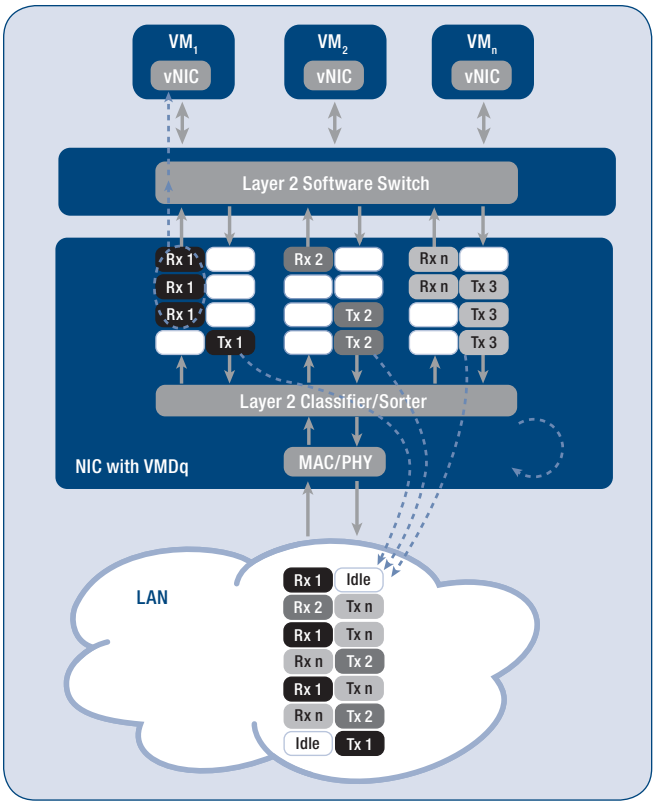

Figure 2. VMDq offloads network I/O sorting to the network silicon, cutting down the processing load on the hypervisor and thereby avoiding I/O bottlenecks and improving throughput.

Transmit Packets

As packets are transmitted from the VMs toward the adapter, the hypervisor transmits data in a round-robin fashion to prevent head-of-line blocking and ensure that each VM is fairly serviced. This guarantees some measure of Quality of Service (QoS) to the VMs.

VMware NetQueue

VMware NetQueue is an interface that is provided in VMware ESX to enable the hypervisor to utilize the VMDq functionality from the Intel® Server Adapter. NetQueue also has the capability to affinitize the VMs to specific CPU cores. By doing so, the data interrupts in the MSI-X vectors for a specific VM are directed to the respective CPU core of that destined VM, as shown in Figure 3. This overall affinitization from the server adapter to the CPU cores of specific VMs improves CPU utilization rates, as a single CPU core is no longer burdened with the data interrupts of all the VMs.

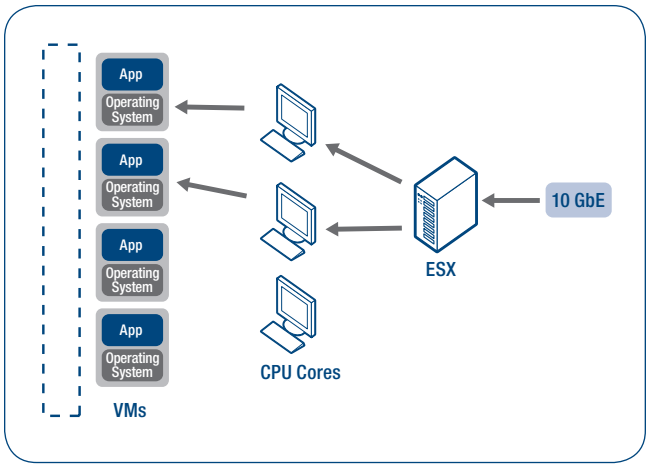

Figure 3. NetQueue improves receive-side networking performance.

The Right Choice of Servers Simplifies IT

To build a next-generation virtualized infrastructure that takes full advantage of intelligent queueing, it is vital to select servers that have been validated to work smoothly with the virtualization technologies that underlie the environment. Dell engineers its PowerEdge servers to ensure optimal interoperability with VMDq and NetQueue technologies. That advance work provides assurance to IT organizations and removes stumbling blocks before they can happen, helping ensure success.

This ecosystem-level enablement helps safeguard the benefits of I/O virtualization to IT managers who might otherwise be daunted by implementing new technologies as they face escalating levels of complexity in day-to-day operations. Many IT organizations struggle to find the resources to innovate at the same time as they must support escalating business needs and growing variety in devices and applications. In response to this widespread need, Dell simplifies IT by removing cost and complexity from the supply chain, adding to the efficiency, power, and flexibility of its hardware.

Dell is currently piloting the *IT Simplification Index*—developed with assistance from three analyst firms—to deliver measurable, tangible results in partnership with its customers in simplifying IT. The index measures efficiency at every stage in the IT lifecycle to identify opportunities to do more with less, assess and improve IT staff efficiency, and make IT more responsive and nimble.

Virtualization solutions based on VMDq and VMware NetQueue are fully validated on Dell PowerEdge servers,[3] providing the power that enterprise environments need for large-scale virtualized workloads and helping to ensure maximum flexibility when implementing this technology to improve network I/O. The 10 Gigabit Ethernet ecosystem enabled by Intel, VMware, and Dell provides the basis for virtualized I/O based on intelligent queueing that eliminates IT headaches, even as it drives up performance.

VMDq Use Case Scenario Demonstrates Performance Benefits

Intel and VMware have collaborated to develop and improve the queueing technology in a virtualized environment. Intel provided its VMDq technology for sorting data packets in the network silicon, which lightens the burden for the hypervisor and improves throughput. VMware improved the hypervisor switch layer, to not only direct the data to the respective destination VM, but also to target interrupts to respective CPU cores and the appropriate VM. With this combined queueing technology implementation in a virtualized environment, the throughput more than doubled. VMware NetQueue and Intel's VMDq combine to efficiently share NICs, increase switching performance with hardware acceleration, enable multiple VMs to be assigned to each port, and allow customers to deploy more applications by reducing CPU utilization.

"For Data Center Managers seeking high performance virtualized I/O in the 10 Gigabit Ethernet era, VMware's network multi-queue technology (NetQueue) and Intel's VMDq provide a highly effective combination for VMware Infrastructure deployments. Together they significantly increase network I/O throughput, switching performance and reduced CPU utilization to allow customers to deploy more applications."

Brian Byun
Vice President of Global Partners and Solutions
VMware

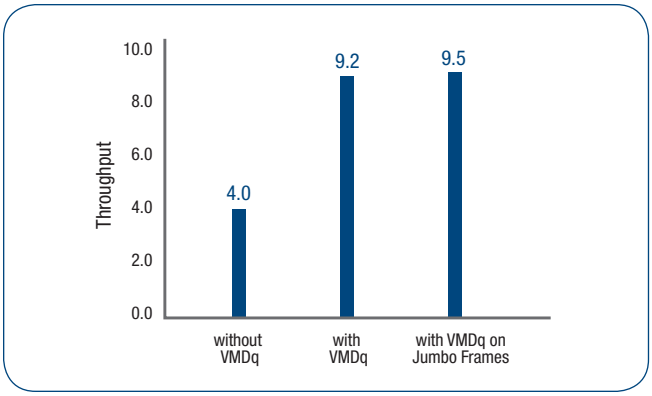

Figure 4. VMDq significantly improves network I/O throughput.

In this specific use case scenario, the configuration included an Intel® Xeon® processor-based server running Windows Server* 2003 with four VMs and an Intel® 82598 10 Gigabit Ethernet Controller running VMware ESX 3.5 Update 1. Using the IxChariot* benchmarking application, receive-only throughput without VMDq was 4.0 Gbps; with VMDq, the throughput more than doubled to 9.2 Gbps. These readings were performed with the standard frame size of 1500 bytes. With 9000-byte Jumbo Frames configured and VMDq enabled, the throughput increased to 9.5 Gbps. VMware supports VMDq on the Intel 82598 10 Gigabit Ethernet Controller in VMware ESX 3.5 Update 1. Results are summarized in Figure 4.

"With excellent engineering and service, as well as proficiency in all facets of supply chains, Dell's OEM division provided us with no-excuses dependability."

Aidymar Bigio
Manufacturing and Operations Manager
Google

Summary

As processor performance continues to increase, IT organizations realize opportunities for higher server-consolidation ratios through virtualization. Growing virtualized workloads mean that customers must now address the associated impacts to I/O. The 10 Gigabit Ethernet virtualized I/O ecosystem based on products and technologies from Intel, VMware, and Dell provides the means to eliminate I/O bottlenecks, safeguarding the scalability and performance of the overall network infrastructure.

VMDq offloads the data packet sorting overhead from the virtual switch in the hypervisor to the Ethernet controller, and VMware NetQueue provides the virtualization software necessary to put this technology to work. Dell PowerEdge servers are pre-validated3 to ensure optimal results from this I/O virtualization, reducing the complexity of IT and making the advantages of the technology more accessible to organizations of all sizes.

Together, these ecosystem components provide data-packet sorting in the network silicon plus individual queues for each VM, freeing more CPU cycles for application processing instead of network I/O processing. That lets your environment work smarter, for better IT results across the board.

For More Information

VMDq is supported in the Intel® 82575 Gigabit Ethernet Controller, Intel® 82576 Gigabit Ethernet Controller, and Intel® 82598 10 Gigabit Ethernet Controller, with appropriate hypervisor enabling. Visit **www.intel.com/go/vtc** for details.

VMware NetQueue is supported in VMware ESX version 3.5 Update 1 and later. Visit **www.vmware.com** for details.

Dell PowerEdge servers help reduce network complexity and maximize your IT resources. Visit **www.DellintelVT.com** for details.

vmware® intel® DELL

[1]Intel® Virtualization Technology requires a computer system with an enabled Intel® processor, BIOS, virtual machine monitor (VMM) and, for some uses, certain platform software enabled for it. Functionality, performance, or other benefits will vary depending on hardware and software configurations and may require a BIOS update. Software applications may not be compatible with all operating systems. Please check with your application vendor.

[2]Intel's Virtual Machine Device queues (VMDq) technology requires a VM vendor's operating system that supports VMDq.

[3]Not all Dell servers are compatible with all adapters. Please check with your Dell sales representative for more information on which servers are supported for the selected adapter.

1008/MBR/PMS/PDF

Order Number: 319823-002US