Intelligent Queueing Technologies for Virtualization

An Intel-VMware Perspective: Enhanced Network Performance in Virtualized Servers

Introduction

As server virtualization continues to grow in IT departments – from small businesses to enterprises – virtualization technologies continue to evolve, improving system throughput for virtual machines and enhancing performance in virtual environments. Virtual Machine Device Queues (VMDq) is another breakthrough technology from Intel that helps offload network I/O data processing from the hypervisor software to the network silicon. When combined with VMware's network multi-queue technology (NetQueue), VMDq improves I/O throughput for faster, more efficient networking.

Table of Contents

Bigger Burdens for 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 increasingly depend on 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.

In virtual environments today, the hypervisor manages network I/O activities. With more virtual machines (VMs) and increased traffic through the platform, the hypervisor requires more CPU cycles to sort data packets and route them to the correct VM (Figure 1), reducing CPU capacity available for applications. Intel's VMDq is a breakthrough technology that reduces the burden on the hypervisor while improving network I/O performance through the virtualized platform.

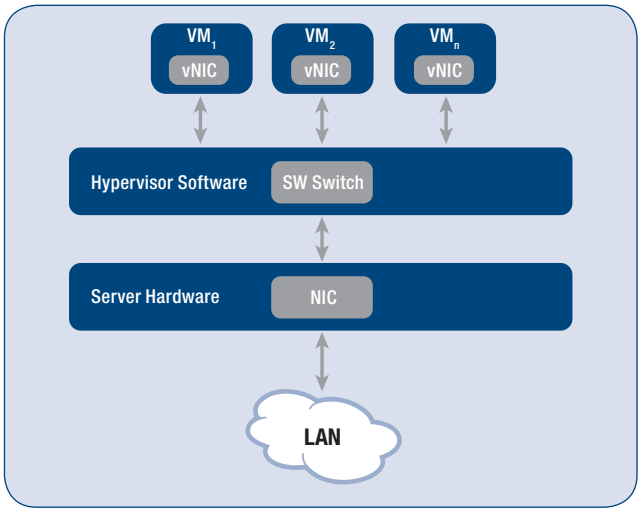

Figure 1. In virtual environments today, the hypervisor manages network I/O.

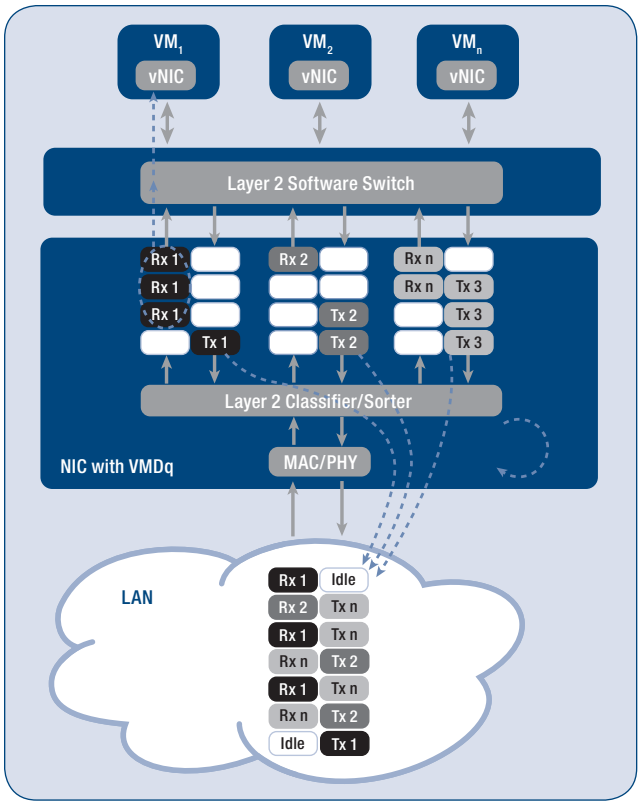

Figure 2. VMDq offloads network I/O management to the network silicon.

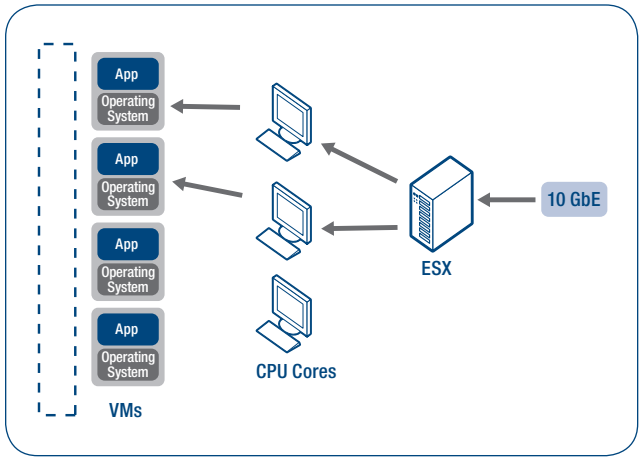

Figure 3. NetQueue improves receive-side networking performance.

Queueing Technology Overview

Intel® Virtualization Technology[1] (Intel® VT) is a set of hardware enhancements that help hypervisor providers develop simpler and more robust virtualization software, plus accelerate system and application solutions in virtual environments. Intel® VT for Connectivity is the portion of Intel VT designed to improve network I/O in virtualized servers. VMDq is part of Intel VT for Connectivity, geared towards improving networking performance and reducing CPU utilization.

VMDq is a network silicon-level technology that offloads network I/O management burden from the hypervisor. Multiple queues and sorting intelligence in the silicon support enhanced network traffic flow in the virtual environment, freeing processor cycles for application work (Figure 2). This improves efficiency in data transactions toward the destined VM, and increases overall system performance.

VMware NetQueue is a performance technology in VMware ESX that significantly improves performance in 10 Gigabit Ethernet virtualized environments. NetQueue provides a network adapter with multiple receive queues that allow data interrupt processing to be affinitized to the CPU cores associated with individual VMs, improving receive-side networking performance. These receive queues can also be assigned to each virtual NIC, mapped to guest memory to avoid a copy, and steer interrupts to idle or optimal processor cores.

Receiving 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.

Transmitting Packets

As packets are transmitted from the VMs towards the adapters, the hypervisor layer places the transmit data packets in their respective queues. To prevent head-of-line blocking and ensure each queue is fairly serviced, the network controller transmits queued packets to the wire in a round-robin fashion, thereby guaranteeing some measure of Quality of Service (QoS) to the VMs.

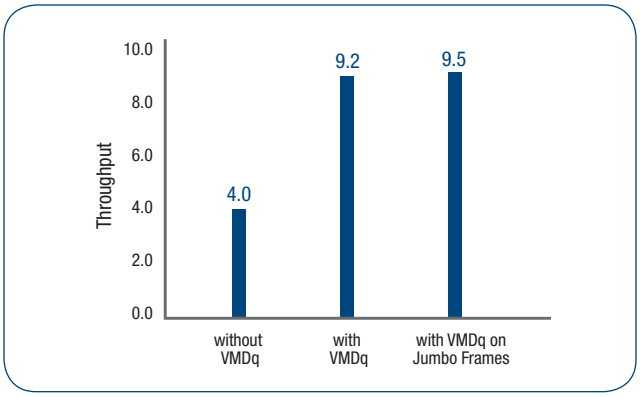

Figure 4. VMDq significantly improves network I/O throughput.

VMDq Performance Use Case Scenario

Intel and VMware have collaborated to develop and improve the queuing 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. VMware improved the hypervisor switch layer, to not only direct the data to the respective destined VM, but also target interrupts to respective CPU cores and their respective destined VM. With this combined queuing technology implementation in a virtualized environment, the throughput more than doubled with a noticeable improvement in CPU utilization. VMware NetQueue and 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.

In this specific use case scenario, the configuration included a Quad-Core Intel® Xeon® processor-based server running Windows* Server 2003 with four VMs and an Intel® 82598 10 Gigabit Ethernet Controller running on 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 with the standard frame size of 1500 bytes. With 9000-byte Jumbo Frames configured and VMDq enabled, the throughput was 9.5 Gbps. VMware supports VMDq on the Intel 82598 10 Gigabit Ethernet Controller in VMware ESX 3.5 Update 1.

Summary

More processing power provides opportunities for greater consolidation in IT data centers; however, the impact to I/O cannot be forgotten. VMDq offloads the data packet sorting overhead from the hypervisor switch to hardware in the network silicon. Data packet sorting in the network silicon, plus individual queues for each VM, free more CPU cycles for application processing instead of network I/O processing. In a benchmark study, the addition of VMDq to Intel network silicon more than doubled the throughput on a virtualized platform.

How to Get VMDq

This feature is supported in Intel® 82575 Gigabit Ethernet Controller and Intel 82598 10 Gigabit Ethernet Controller, and needs appropriate hypervisor enabling.

How to Get VMware NetQueue

VMware NetQueue is supported in version 3.5 and later of VMware ESX.

Visit **www.intel.com/network** or **www.vmware.com** for details.

vmware®

intel®

[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.

Printed in USA 0408/BY/OCG/HOP/500 Order Number: 319823-001US