TECHNICAL WHITE PAPER

Advanced Virtualization I/O Queuing Technologies

An Intel-Microsoft Perspective

TABLE OF CONTENTS

Enhanced Queuing Technologies for Improved Network Performance in Virtualized Servers

With the increase in the adoption of server virtualization and consolidation in data centers, there has been a continuous evolution of virtualization technologies in server components to improve the overall performance of the system and achieve the Service Level Objectives (SLOs) as defined. To improve the I/O performance of these virtualized servers, Intel has developed a breakthrough technology, Virtual Machine Device Queues (VMDq). VMDq helps off load network I/O data processing from the hypervisor software to the network silicon. Microsoft has developed VMQ technology, available in Windows Server® 2008 R2 Hyper-V.™ These two technologies work together to improve I/O throughput for faster, more efficient networking.

Bigger Burdens for Hypervisors

Deploying virtualized environments is a growing practice among IT departments in order to consolidate server workloads and reduce data center footprints. Today's more powerful servers allow for greater virtual machine (VM) density per virtualized server than ever before, yet this consolidation does not necessarily mean more efficient network throughput in the virtual environment. In fact, it can have a significant impact on system and application performance as workloads increasingly depend on network I/O. A balance between system performance and networking capabilities is required to achieve optimal application services from consolidation.

In the Hyper-V environment today, the virtual switch in the Management OS filters data based on VLAN tags and MAC addresses, copies the data, and then routes the data to the respective VMs via the VM Bus as shown in Figure 1. The overhead associated with the virtual switch and the data copies impact a system's overall network I/O performance, including both CPU utilization and throughput.

The solution provided by Microsoft and Intel with their complementary queuing technologies, VMQ and VMDq, reduces this impact to improve system performance.

Queuing Technology Overview

Intel® Virtualization Technology[1] (Intel® VT) refers to the hardware assists for virtualization that Intel offers across its server platforms (CPU, Chipset, I/O) to provide improved system performance, security, efficiency, and a more powerful virtualization solution. Intel® VT for Connectivity is the portion of Intel VT designed to improve network I/O in virtualized servers and includes VMDq.

VMDq is a network silicon-level technology that off loads the network I/O management burden from the hypervisor to the Ethernet Controller. 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.

VMQ is Microsoft's Hyper-V queuing technology that makes use of the VMDq capabilities of the Intel Ethernet controller to enable data packets to be delivered to the VM with minimal handling in software. The Shared Memory feature allows the data packets to DMA directly into the VM's memory, thereby avoiding a copy between the memory of the Management OS and the VM's memory.

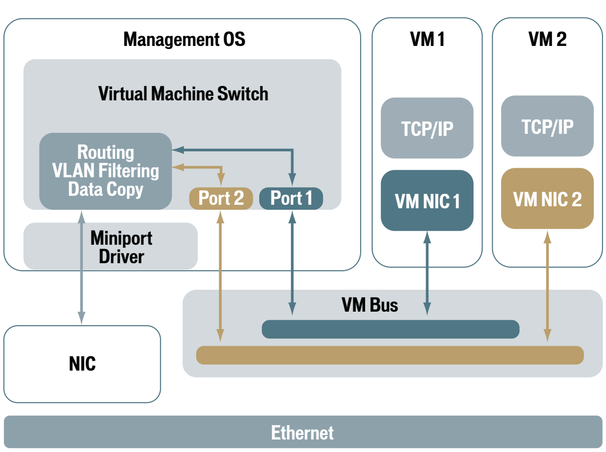

Figure 1. In virtual environments today, the hypervisor's virtual switch manages network I/O. All data passes through the virtual switch in the Management OS.

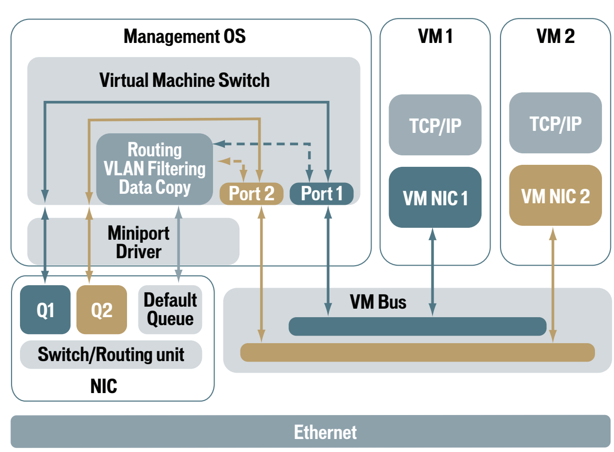

Figure 2. With VMQ and VMDq technologies, data packets are delivered directly to the VM without data copy and bypassing the virtual switch.

Receiving Packets

As data packets arrive at the network adapter, a Layer 2 classifier/sorter in the Ethernet controller sorts and determines which VM each packet is destined for based on MAC addresses and/or VLAN tags that the Hyper-V virtual switch has programmed into the controller. The data packets are sorted into multiple queues in the controller by VMDq, and then VMQ's Shared Memory feature DMAs them directly into the VMs' memory. Doing so removes the overhead associated with the filtering and data copy overhead from the virtual switch in the Management OS. Thus, VMDq improves platform efficiency for handling receive-side network I/O and lowers CPU utilization for application processing in a Hyper-V environment.

Transmitting Packets

As packets are transmitted from the VMs towards the adapters, the hypervisor places the 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.

Performance Use Case Scenario

Together, Intel and Microsoft's queuing technologies improve overall network I/O performance in a virtualized environment. With this combined queuing technology implementation in a virtualized environment, lab tests demonstrated significant throughput improvement. Microsoft's VMQ and Intel's VMDq combine to efficiently share network I/O, increase switching performance with hardware acceleration, and allow customers to deploy more applications.

Microsoft's VMQ and Intel's VMDq combine to efficiently share network I/O, increase switching performance with hardware acceleration, and allow customers to deploy more applications.

In this specific use case scenario, the configuration included a dual-socket Intel® Xeon® processor-based server, for a total of eight cores, running Windows Server 2008 Hyper-V R2 Beta, and an Intel® 82598 10 Gigabit Ethernet Controller. Using the ntttcp application, receive-only throughput without VMDq was 5.4 Gbps for 4 VMs; with VMDq, the throughput went up to more than 9.3 Gbps.[3] These readings were with the standard frame size of 1500 bytes.

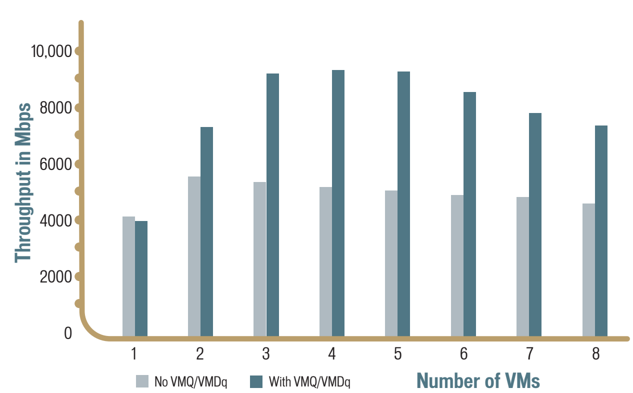

Figure 3. Together, VMQ and VMDq significantly improve network I/O throughput.[2]

Summary

As server processing power increases and drives greater VM density on physical servers, the impact on network I/O will need to be addressed. Together, VMQ and VMDq improve virtualized server I/O by off loading the data packet sorting overhead from the hypervisor virtual switch to the Ethernet controller. Data packet sorting in the Ethernet controller, plus individual queues for each VM and the DMA of the data packets directly to the VM's memory, make more CPU cycles available for application processing instead of network I/O processing and improve server throughput. Customers can realize these benefits by deploying Intel® Ethernet Server Adapters and Windows Server 2008 R2 Hyper-V.

How to Get VMDq

VMDq is supported in the following products:

- 1GbE (Intel® 82575 Gigabit Ethernet Controller, Intel® 82576 Gigabit Ethernet Controller)
- 10GbE (Intel® 82598 10 Gigabit Ethernet Controller, Intel® 82599 10 Gigabit Ethernet Controller)

How to Get Microsoft VMQ

VMQ is included in the Windows Server 2008 R2 with the Hyper-V role enabled.

Visit **www.intel.com/network, www.intel.com/go/vtc** or **www.microsoft.com** for details.

[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] This is throughput measured by TCP. At 4 VMs, this effectively represents line rate when IP and Ethernet overhead is included.

[3] Microsoft internal measurement. (April 2009) ntttcp benchmark. Windows Server 2008 R2 Beta. Intel® Xeon® processor E5460, 3.16 GHz, 8 MB L2 cache, 1333 MHz system bus, 24 GB memory.

Printed in USA 0509/TAR/OCG/XX/PDF ♻ Please Recycle 321968-001US