



White Paper
Intel® Virtualization Technology

Virtual Machine Device Queues

An Integral Part of Intel® Virtualization
Technology for Connectivity that Delivers
Enhanced Network Performance

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 virtual machine monitor software to the network silicon.

Shefali Chinni
Product Marketing Manager
for I/O Virtualization

Radhakrishna Hiremane
Product Marketing Engineer
Server Platforms Group Marketing

Contents

Bigger Burdens for Hypervisors	2
Virtual Machine Device Queues (VMDq) Overview	2
Receiving Packets	3
Transmitting Packets	3
VMDq Performance Use Case Scenario	3
Summary	3
How to get VMDq?	4

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. This practice, however, can have a significant impact on system and application performance as workloads increasingly depend on network I/O. While IT managers add greater processing power and reduce 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 Virtual Machine Device Queues (VMDq) is a breakthrough technology that reduces the burden on the hypervisor while improving network I/O performance through the virtualized platform.

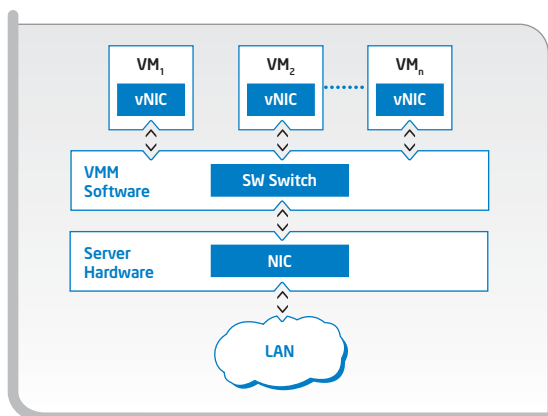


Figure 1. In virtual environments today, the hypervisor manages network I/O.

Virtual Machine Device Queues (VMDq) Overview

Intel® Virtualization Technology (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. Virtual Machine Device Queues (VMDq) is part of Intel® VT for Connectivity, geared towards improving networking performance and reducing CPU utilization.

Virtual Machine Device Queues (VMDq) is a 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.

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 increases CPU utilization for application processing.

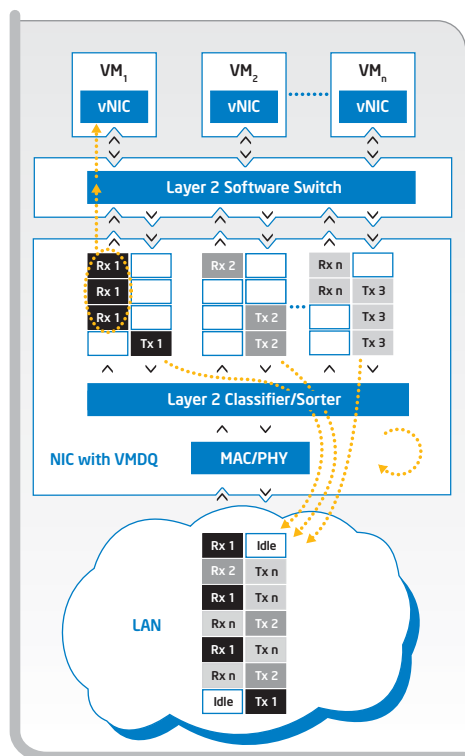


Figure 2. VMDq offloads network I/O management to the network silicon.

Transmitting Packets

As packets are transmitted from the virtual machines 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 via queues.

VMDq Performance Use Case Scenario

Intel and VMware have partnered together 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 upon the hypervisor switch layer, to direct not only 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.

In this specific use case scenario, the configuration included a Quad-Core Intel® Xeon® processor-based server running Windows* 2003 with 4 VMs, Intel® 82598 10 Gigabit Ethernet Controller running on an ESX development build. Without VMDq, the throughput 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 Jumbo Frames, the throughput was 9.5 Gbps. VMware plans to support VMDq on Intel® 82598 10 Gigabit Ethernet Controller in a future version of ESX.

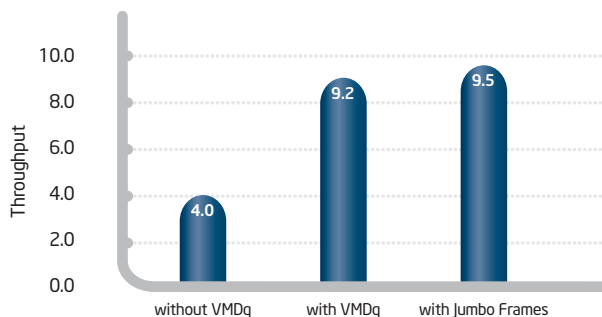



Figure 3. VMDq significantly improves network I/O throughput.

Summary

More processing power provides opportunity for greater consolidation in IT data centers; however, the impact to I/O cannot be forgotten. Virtual Machine Device Queues offload 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.



www.intel.com

How to get VMDq?

This feature is supported in Intel® 82575 Gigabit Ethernet Controller and Intel® 82598 10 Gigabit Ethernet Controller, and requires virtualization software enabling.

Visit: www.intel.com/network for details