

By Ravikanth Chaganti
Barun Chaudhary

VIRTUAL MACHINE QUEUES: ENHANCED THROUGHPUT FOR MICROSOFT WINDOWS SERVER 2008 R2 HYPER-V

The virtual machine queue feature introduced in the Microsoft® Windows Server® 2008 R2 Hyper-V™ platform is designed to offload virtual network processing to physical adapter hardware. Enabling this feature on Dell™ PowerEdge™ servers with supported Intel® adapters can help significantly increase overall throughput while reducing the network processing burden on host servers.

Although consolidating application workloads onto virtual machines (VMs) can help significantly increase server utilization, it also introduces its own specific challenges, including contention for available network bandwidth and other resources. The Microsoft Windows Server 2008 R2 OS introduces several enhancements to Hyper-V designed to address these challenges, one of which is the VM queue (VMQ) feature. This functionality requires support from the underlying network adapter hardware, but as Dell testing shows, enabling the feature can dramatically increase overall throughput for network I/O-intensive workloads while helping reduce the network processing burden on host servers.

VMQ ARCHITECTURE AND FUNCTIONALITY

In previous versions of Hyper-V, the virtual-network-switch implementation processed network traffic between VMs and physical network interface cards (NICs). This processing included classifying and filtering packets based on Media Access Control (MAC) addresses or virtual LAN (VLAN) identifiers as well as routing these packets to the appropriate VM adapter. Because the number of packets being routed by the host server typically increases when VM density rises, so does host processor utilization.

The VMQ feature introduced in Windows Server 2008 R2 Hyper-V is designed to overcome this challenge and enable optimal use of hardware resources. This feature is a silicon-level implementation that works in conjunction with hardware support and includes the intelligence for multiple queue formation (as supported by the specific adapter) and the sorting of packets in the physical NIC. This approach offloads packet sorting and routing from the host processor, helping save processor cycles for other applications. Organizations can deploy this technology using supported Dell PowerEdge platforms with Gigabit Ethernet (GbE) Intel VT quad-port server adapters as well as 10 Gigabit Ethernet (10GbE) Intel XF SR, Intel AT, and Intel DAC server adapters; other upcoming 10GbE Intel adapters are expected to support the feature as well.

Figure 1 shows the VMQ architecture. Without VMQ enabled, the hypervisor must sort and route the packets to the destination VMs, and the incoming packets are copied twice while transmitting from the NIC buffer to the application buffer. With VMQ enabled, multiple queues are formed, with each queue corresponding to a VM; there is also a default queue that routes multicast packets, broadcast packets, and packets without a specified queue. When data packets arrive at the physical NIC, the NIC sorts the packets based on

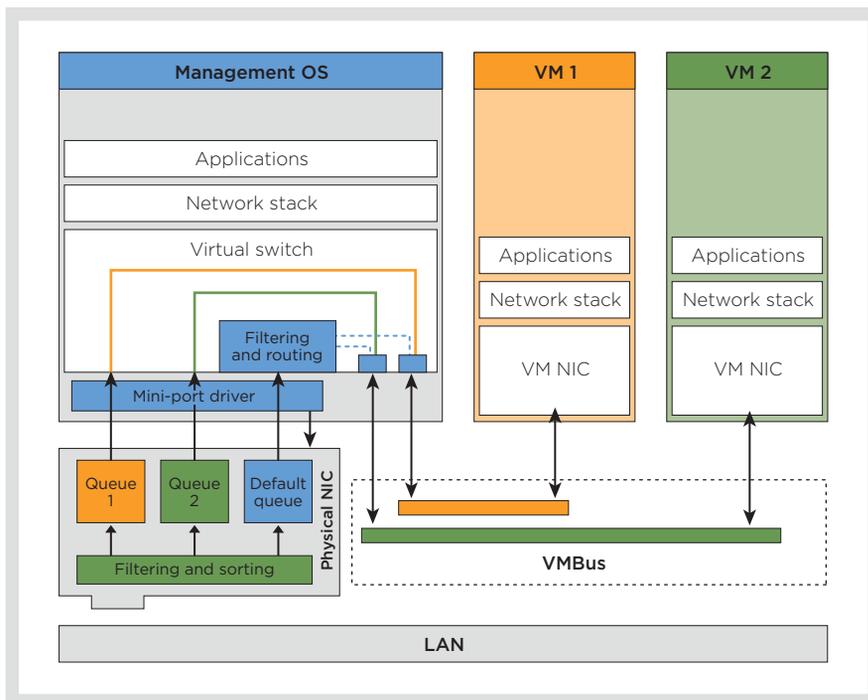


Figure 1. Microsoft Windows Server 2008 R2 Hyper-V VMQ architecture

the MAC addresses and VLAN tags of the destination VM and places them in the appropriate receive queues. The virtual network switch then takes the corresponding packets in the queue to the VMs.¹

VMQ CONFIGURATION

The Intel network adapter drivers for Dell servers disable VMQ by default. To enable VMQ, administrators should enable VMQ support on the physical network adapter and on the virtual switch. There is no configuration required inside the guest OS to enable this feature.

Enabling VMQ support for the virtual switch first requires modifying the registry:

- For GbE adapters, administrators can perform this step by adding a `BelowTenGigVmqEnabled` entry under `HKLM\SYSTEM\CurrentControlSet\Services\VMSMP\Parameters` and setting its value to 1.
- For 10GbE adapters, they would add a `TenGigVmqEnabled` entry in the same location, also with its value set to 1.

Next, to enable the VMQ feature on the physical network adapter, administrators can launch the Device Manager tool, expand the “Network adapters” section, right-click on the adapter that supports the VMQ feature, and select “Properties.” In the properties window, they can then select the Virtual Machine Queues property, set the value to “Enabled” (see Figure 2), click the OK button, and close Device Manager.

Administrators should note that they must perform these steps in the order described.

When using Microsoft System Center Virtual Machine Manager (SCVMM) 2008 R2 to manage Hyper-V VMs, administrators can enable or disable the VMQ feature for a VM network adapter by selecting or unselecting the Enable Virtual Network Optimizations check box in the VM properties window. This option is available only when using SCVMM 2008 R2 with Windows Server 2008 R2 Hyper-V VMs.

NETWORK THROUGHPUT FOR VIRTUALIZED FILE SERVERS

Network I/O-intensive workloads such as file servers can directly benefit from using the VMQ feature on Windows Server 2008 R2 Hyper-V host servers. To demonstrate this enhancement, in August 2009 Dell engineers performed several tests using custom scripts to stress a file server with and without the VMQ feature enabled, and measured the resulting network throughput and file copy times.

The test environment consisted of one file server and three client servers. The file server was a Dell PowerEdge R610 configured with one quad-core Intel Xeon® E5530 processor at 2.4 GHz, 8 GB of RAM, a 10GbE Intel AT server adapter, four GbE Broadcom LAN on Motherboards (LOMs), and the Windows Server 2008 R2 Datacenter OS. The three clients were PowerEdge SC1435 servers, each configured with one quad-core AMD Opteron™ 2350 processor at 2.0 GHz, 4 GB of RAM, two GbE Broadcom LOMs, and the 32-bit version of the Microsoft Windows Vista® OS with Service Pack 1 (SP1).

The custom scripts were designed to copy large numbers of mixed-size files (varying from a few kilobytes to approximately 6 GB) from the clients to the VMs. The VMs were stored on an external Dell PowerVault™ MD3000i Internet SCSI

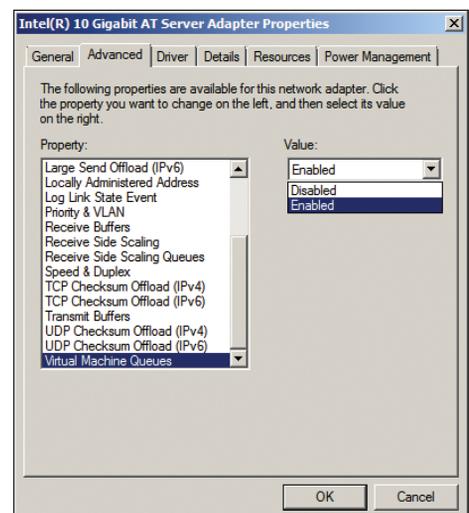


Figure 2. Adapter properties in the Microsoft Windows Server 2008 R2 Device Manager tool

¹For more information on the VMQ feature in Microsoft Hyper-V, visit msdn.microsoft.com/en-us/library/dd568132.aspx. For more information on this feature in Intel server adapters, visit www.intel.com/network/connectivity/vtc_vmq.htm.

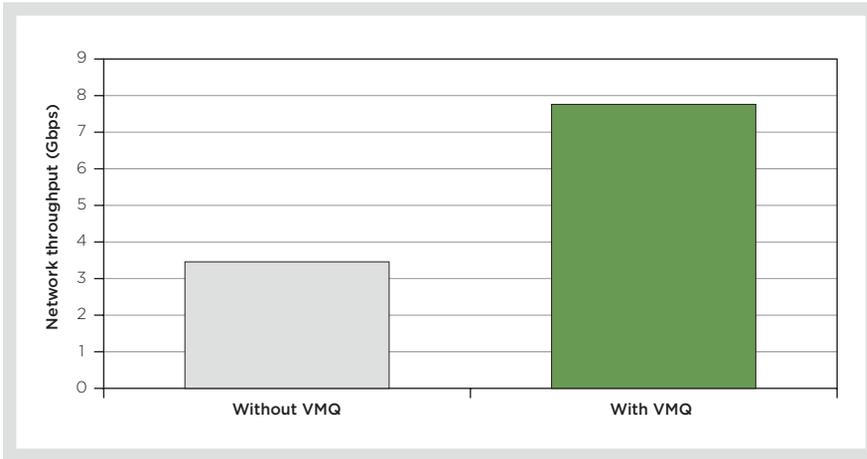


Figure 3. Network throughput with and without the VMQ feature enabled

(iSCSI) storage array, and connected to a virtual switch configured on the 10GbE Intel AT adapter; the clients connected to GbE ports on the same switch as the VMs. The iSCSI logical units (LUNs) were directly mapped into the guest OS using the Microsoft iSCSI Software Initiator to help ensure that the disk I/O load translated to a network I/O load, and therefore resulted in an increased load on the virtual network.

These tests showed that enabling the VMQ feature reduced the average number of hardware interrupts to the management OS—translating to reduced processor utilization on the host server—and helped increase overall network throughput. Figure 3 shows the network throughput with and without the VMQ feature enabled:

enabling the VMQ feature with the 10GbE network adapter more than doubled the throughput, to approximately 7.8 Gbps.

The increase in overall network throughput with the VMQ feature enabled also led to reduced copy times between the test clients and the VMs. Figure 4 shows these copy times with and without the VMQ feature enabled: enabling the feature reduced the times by more than 60 percent. These tests also demonstrate the enhanced file server application throughput when using VMQ-enabled virtual switch ports.

ENHANCED PERFORMANCE FOR VIRTUALIZED ENVIRONMENTS

Dell engineering teams worked with Microsoft and Intel from the early stages

of Windows Server 2008 R2 development to validate its enhanced Hyper-V functionality, and a broad selection of Dell PowerEdge server models support the GbE and 10GbE Intel adapters with the VMQ feature. Although GbE network adapters do support VMQs, best practices recommend using 10GbE network adapters to help maximize the advantages of the feature. For environments with supported hardware, enabling the VMQ feature can help significantly increase throughput for network I/O-intensive workloads such as file services running inside a VM while also helping to reduce the processing burden on host servers.

Ravikanth Chaganti is a lead engineer on the Windows Server OS team at the Dell India R&D Center in Bangalore. He has been with Dell for the past six years and has worked on multiple Microsoft Windows® OS releases. His current interests include Windows Server virtualization and OS performance.

Barun Chaudhary is a software engineer on the Windows Server OS team at the Dell India R&D Center in Bangalore. His current interests include Windows Server virtualization and networking technology. He has a B.Tech. in Electrical Engineering from Motilal Nehru National Institute of Technology, Allahabad.



Figure 4. File copy times with and without the VMQ feature enabled

MORE
ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Microsoft Windows Server 2008 R2:
DELL.COM/WindowsServer2008
www.microsoft.com/windowsserver2008

Microsoft Hyper-V:
DELL.COM/HyperV
www.microsoft.com/hyperv