

Greater Manageability and Throughput

Cisco® Nexus™ 1000V Series Switches and Intel's Virtual Machine Device Queues



Server virtualization has numerous benefits in terms of reduced equipment and power costs, space savings, and management. However, it also introduces new challenges, including greater overhead due to increased I/O throughput and limitations in server access network elements such as services and policies.

Cisco and Intel provide innovative solutions in order to address these challenges:

- The Cisco® Nexus™ 1000V Series Switch is a pure software implementation of a Cisco Nexus switch. It resides on a server and integrates with the VMware ESX* hypervisor to deliver virtual machine-aware network services in a virtualized environment.
- Intel's Virtual Machine Device Queues (VMDq) technology addresses the increasing amount of traffic on virtualized servers by offloading data sorting from the hypervisor's virtual switch (vSwitch) to reduce I/O overhead and deliver near-native I/O throughput.

Cisco Nexus 1000V Series Switches

Cisco Nexus 1000V Series Switches are virtual machine access switches for VMware vSphere* environments running the Cisco® NX-OS Software operating system. Operating inside the VMware ESX hypervisor, the Cisco Nexus 1000V Series Switch provides:

- Policy-based virtual machine connectivity
- Mobile virtual machine security and network policy
- Non-disruptive operational model for your server virtualization and networking teams

The Cisco Nexus 1000V Series Switch has two major components: the virtual Ethernet module (VEM) and the external virtual supervisor module (VSM).

The Cisco Nexus 1000V Series VEM runs as part of the VMware ESX or ESXi kernel and replaces the VMware vSwitch. The VEM applies network policies to traffic bound for VMs on the server.

The Cisco Nexus 1000V Series VSM hosts defined policies and controls multiple VEMs as one logical modular switch.

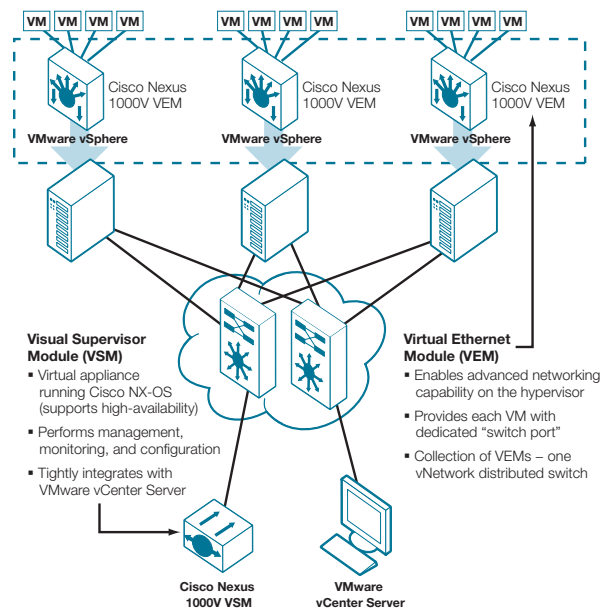


Figure 1. Cisco® Nexus™ 1000V Series Switch architecture

Developed in close collaboration with VMware, the Cisco Nexus 1000V Series is certified by VMware to be compatible with VMware vSphere, vCenter,* ESX, and ESXi, and with many other vSphere features.

Virtual Machine Device Queues

Virtual Machine Device Queues (VMDq) is a breakthrough network silicon-level technology from Intel that offloads the network I/O management burden from the vSwitch, improving I/O throughput in virtualized servers. Introduced in 2007, VMDq has been deployed by many customers to improve networking performance in VMware ESX environments.

VMDq reduces I/O overhead on the hypervisor by performing data sorting in the Intel® Ethernet controller, rather than the vSwitch. As data packets arrive at the network adapter, a Layer 2 classifier/sorter in the Ethernet controller sorts and

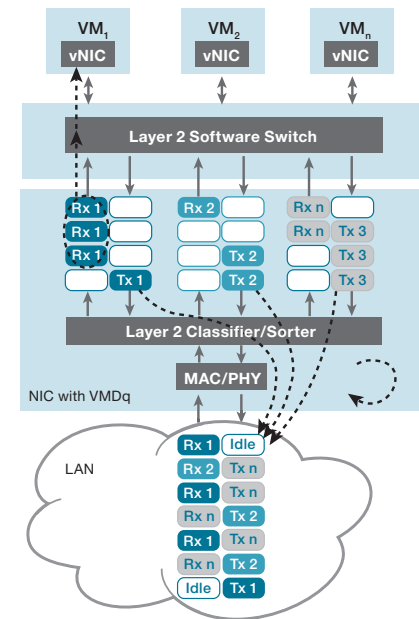


Figure 2. VMDq offloads data sorting to the Intel® Ethernet controller

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 vSwitch then routes the packets to the respective VM instead of performing the heavy lifting work of sorting data.

VMDq, available in 1GbE and 10GbE Intel Ethernet controllers and adapters, is a key component of Intel® Virtualization Technology for Connectivity (Intel® VT for Connectivity), a suite of hardware assists designed to improve networking performance and improve system performance. Intel VT for Connectivity is an important element of Intel® Virtualization Technology,¹ a set of platform hardware enhancements that help hypervisor providers develop simpler and more robust virtualization software, plus accelerate system and application solutions in virtual environments.

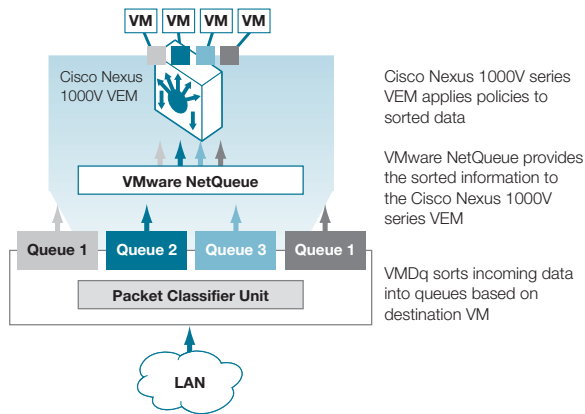


Figure 3. Cisco® Nexus™ 1000V Series Switch and VMDq in a VMware ESX® environment

Working Together

Tight integration with VMware ESX allows the Cisco Nexus 1000V Switch and Intel's VMDq to work together seamlessly to deliver near-native throughput and VM-aware networking. When incoming data packets arrive at the Intel Ethernet port, VMDq groups them into queues. NetQueue, VMware's network multiqueue technology, then forwards the queued packets to the Cisco Nexus 1000V Switch. Upon receipt of the data packets, the Cisco Nexus 1000V Switch identifies the destination address, and deploys the policies associated with the destination VM, before forwarding the packets to the VMs.

Performance

Testing performed by Intel has shown that an Intel® Ethernet 10 Gigabit Server Adapter with VMDq enabled achieves near-native throughput in a VMware ESX-based server with both the ESX vSwitch and the Cisco Nexus 1000V Switch.²




Figure 4. VMDq performance with VMware ESX® vSwitch and Cisco® Nexus™ 1000V Switch

Summary

Today's more powerful servers provide opportunities for greater consolidation through virtualization in IT data centers; however, the impact to server network management and I/O throughput cannot be forgotten. The Cisco Nexus 1000V Series Switches and VMDq from Intel address these issues to help increase data center efficiency and scalability.

With the Cisco Nexus 1000V Series, virtual machines are treated in the same way as physical servers in security policy, monitoring and troubleshooting. Nexus 1000V extends the existing operational model between network and server administrators and enables virtual machines to be true basic building blocks of the data center. These operational efficiencies lead to greater scaling of server virtualization deployments with lower operation expenditures.

VMDq's ability to offload the data packet sorting overhead from the vSwitch to hardware in the Intel Ethernet controller, frees more CPU cycles and improves throughput, allowing for greater application processing and increased VM scalability.

For More Information

For more information about Cisco Nexus 1000V Series Switches, visit www.cisco.com/go/1000v

For more information about Intel Ethernet products, visit www.intel.com/go/Ethernet

For more information about Intel Virtualization Technology for Connectivity, visit www.intel.com/go/vtc

¹Intel® 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.

²Performance claim supported by performance result of a benchmark using NTttcp. Intel internal measurement (June 2009). Production system with two Intel® Xeon® processors 5500 series 2.93 GHz, two 10GbE adapters based on the Intel® 82598 10 Gigabit Ethernet Controller, VMware ESX 4.0. Systems were connected using a Cisco Catalyst® 6503 switch with a 16-port 10GBase-T module.

Copyright © 2009 Intel Corporation. All rights reserved. Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

Copyright © 2009 Cisco Systems, Inc. All rights reserved. Cisco, the Cisco logo, and Nexus are trademarks or registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

*Other names and brands may be claimed as the property of others.