

Intel® Virtualization Technology for Connectivity

Each new generation of processor brings more processing power and new capabilities to server platforms, and today's servers are running more applications and processes simultaneously than ever before. IT departments are deploying 10 Gigabit Ethernet (10GbE) to meet the I/O demands of these servers and also to support converged LAN and SAN networking. 10GbE's increased throughput and support for multiple traffic types gives IT administrators the opportunity to simplify their network infrastructures and increase network flexibility to adapt to changes in demand quickly and effectively.

Intel® Virtualization Technology for Connectivity¹ (Intel® VT-c) is a key feature of many Intel® Ethernet Controllers. With I/O virtualization and Quality of Service (QoS) features designed directly into the controller's silicon, Intel VT-c enables I/O virtualization that transitions the traditional physical network models used in data centers to more efficient virtualized models by providing port partitioning, multiple Rx/Tx queues, and on-controller QoS functionality that can be used in both virtual and non-virtual server deployments.

On-Controller Quality of Service for Virtualized and Non-virtualized Environments

At the core of Intel VT-c are functions and technologies that are integrated into the Intel Ethernet Controller to provide a common QoS feature set that delivers a variety of capabilities that can be used directly by an OS or hypervisor or configured by an IT Administrator to meet a specific need (Figure 1).

Virtual Machine Device Queues (VMDq) and PCI-SIG* Single Root I/O Virtualization (SR-IOV) are two of the optimization technologies used to enable the enhanced I/O virtualization functions of Intel VT-c. These functions help reduce I/O bottlenecks and improve overall server performance by offloading functionality to the Intel Ethernet Controller. They share many of the same QoS features integrated into the controller's silicon, and they both provide native throughput, balanced bandwidth allocation, and improved I/O scalability.

In virtualized servers, these functions help reduce I/O bottlenecks and improve the overall server performance by offloading the data sorting and queuing functionality from the hypervisor to the Intel Ethernet Controller.

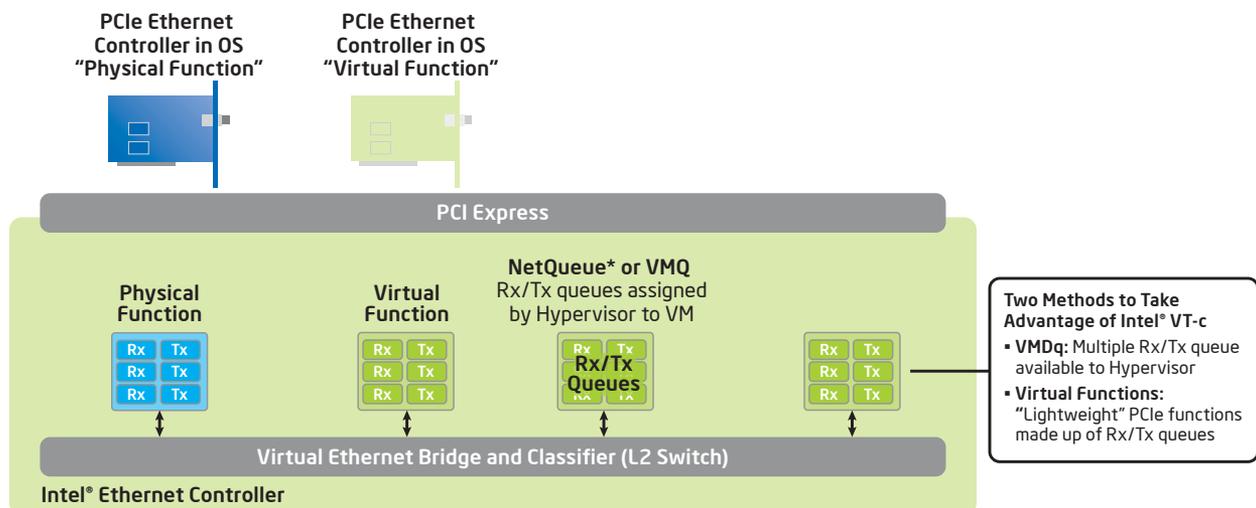


Figure 1. Overview of Intel® Virtualization Technology for Connectivity

In non-virtualized environments, port partitioning provides hardware-based QoS features that enable a physical port to consolidate the traffic of a greater number of physical ports with no loss of functionality and more flexible bandwidth allocation.

Both technologies enable a single Ethernet port to appear as multiple adapters to virtual machines (VMs) by allowing the Intel Ethernet Controller to place data packets directly into individual VM memory stacks using a process called direct memory access (DMA). Each VM's device buffer is assigned a transmit/receive queue pair in the Intel Ethernet Controller, and this pairing between VM and network hardware helps avoid needless packet copies and route lookups in the virtual switch. The result is less data in the host server's buffers and an overall improvement in I/O performance.

Hypervisor-controlled Port Partitioning Using Virtual Machine Device Queues

VMDq works in conjunction with VMware NetQueue* or Microsoft Virtual Machine Queues* (VMQ), in their respective hypervisors, to use the sorting and queuing functionality in the controller for traffic steering and Tx/Rx round-robin scheduling for balanced bandwidth allocation across multiple transmit and receive queues. Using these technologies,

VMDq enables the hypervisor to represent a single network port as multiple ports that are assigned to the VMs, resulting in less data in the host's buffers and an overall performance improvement to I/O operations. When VMDq is enabled, the hypervisor handles queue assignment, delivering the benefits of port partitioning and the on-controller QoS features with little to no administrative overhead by the IT staff (Figure 2).

SR-IOV-enabled Port Partitioning

Many recent Linux* releases have been enabled to partition a single physical Ethernet Controller into multiple virtual interfaces that can be used by local host processes or directly by VMs. With support for the PCI-SIG SR-IOV specification, Intel Ethernet Controllers support this port partitioning, which administrators can use to create multiple isolated networks for use in both bare-metal Linux and virtualized server deployments. As seen in Figure 3, the Linux OS screenshots show the two ports on the Intel® 82576 Gigabit Network Connection available on the server as Ethernet Controllers. After using the max_vfs command to enable seven virtual functions per port, the OS shows the original two ports and 14 additional Ethernet Controllers that are now available for the administrator to assign to traffic to.

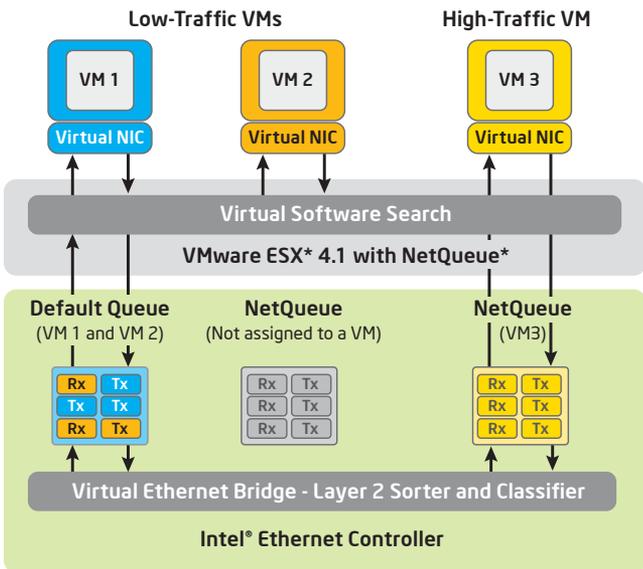


Figure 2. Virtual Machine Device Queues

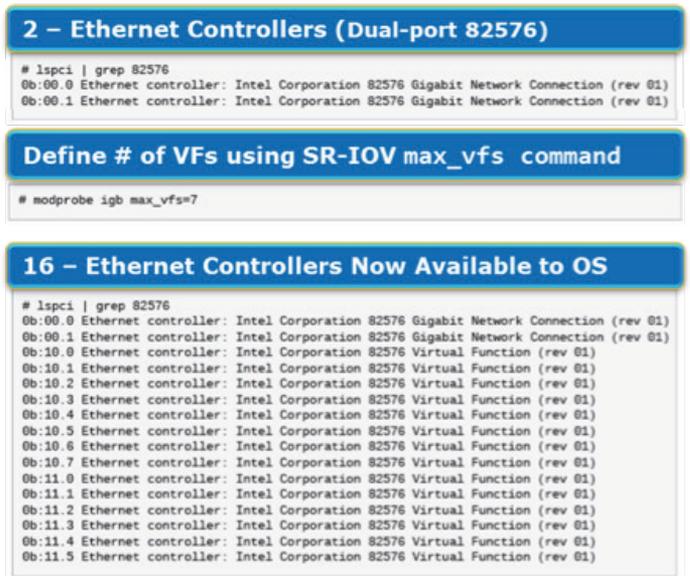


Figure 3. Linux OS screenshots

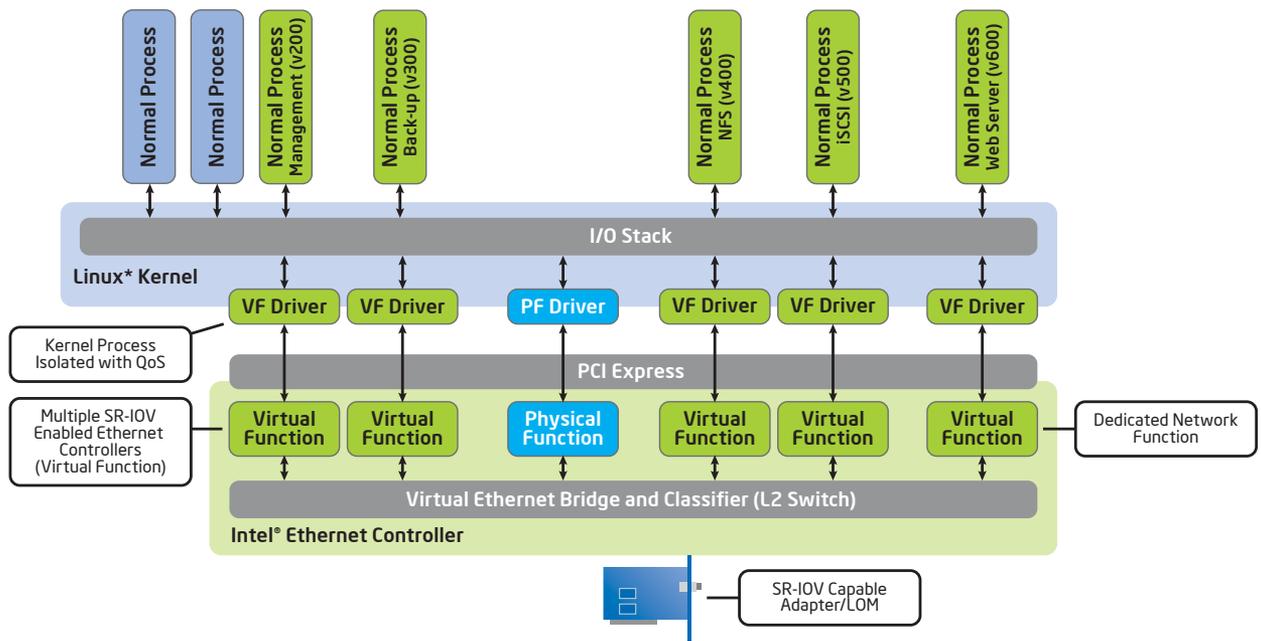


Figure 4. Port Partitioning implemented in a bare-metal Linux* environment

In a bare-metal Linux server, host processes can be assigned to these dedicated network resources to provide traffic isolation and balanced bandwidth allocation. Hardware-based QoS functionality in the Intel Ethernet Controller keeps the network connections available for critical traffic during heavy traffic contention. These isolation and QoS features are critical in 10GbE network deployments, but they can also be used in Gigabit Ethernet (GbE) environments to help ensure that management traffic has access to the host OS (Figure 4).

In a virtualized environment, the direct assignment of a VM to a virtual adapter reduces the CPU overhead seen when using a software-based network bridge or virtual switch by offloading network traffic management to the Intel Ethernet Controller. New capabilities included in Linux drivers allow the configuration of hardware bandwidth throttling and monitoring capabilities, allowing fine tuning of QoS requirements for each virtual adapter whether it is used by the OS or by a VM. Figure 5 shows SR-IOV used in a combined bare-metal Linux and virtualized environment.

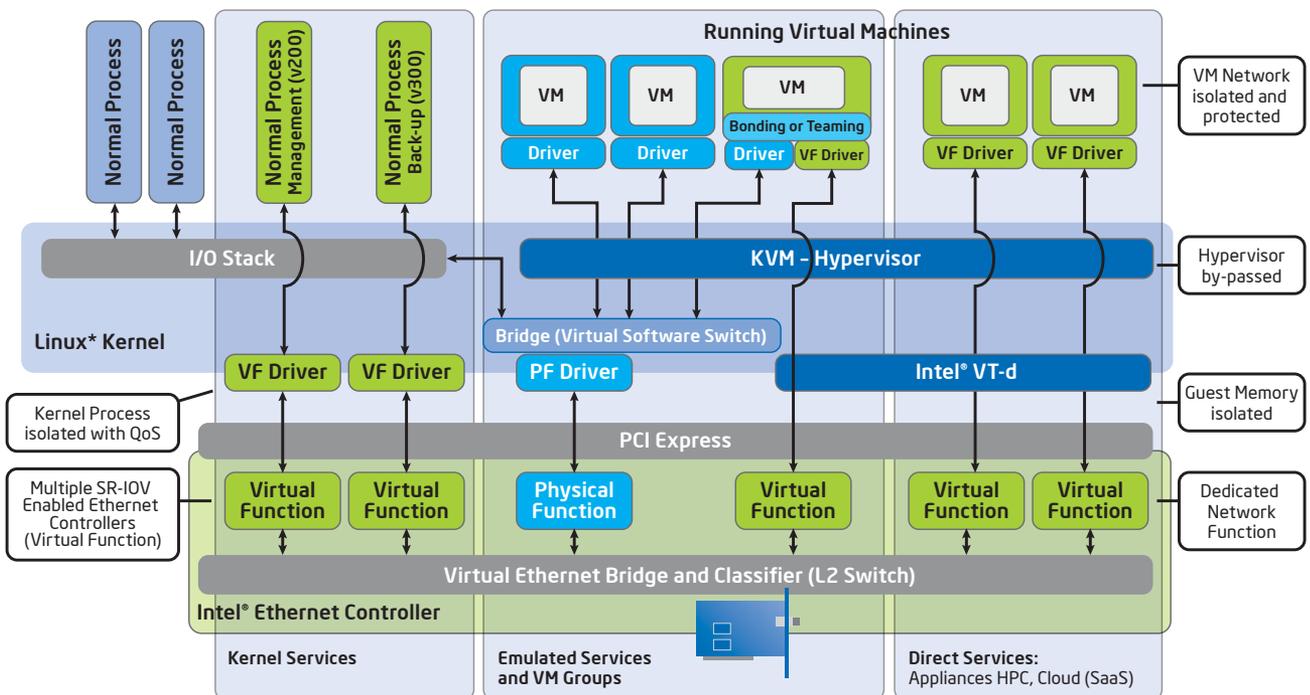


Figure 5. Port Partitioning in a combined bare-metal Linux* and virtualized environment

Conclusion

Growing deployments 10GbE give IT administrators the opportunity to create simpler, more flexible network infrastructures. Integrated into Intel Ethernet Controllers, Intel Virtualization Technology for Connectivity provides functionality that can help these administrators address their various and unique I/O needs for both virtualized and non-virtualized servers.

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.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information. The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order. Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web Site <http://www.intel.com/>.

Copyright © 2011 Intel Corporation. All rights reserved. Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

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

