



Allocating bandwidth dynamically for quality of service delivery

By Kyle Roblyer and Rahul Deshmukh



Next-generation I/O

Dell and QLogic are driving next-generation I/O for optimizing networking performance in virtualized environments. View this demonstration to learn more about how Switch Independent Partitioning on QLogic QLE8262 CNAs enhances network QoS in virtualized data centers.

qrs.ly/582gdco

Virtualization and converged networks bring tremendous benefits and cost savings, but may create I/O bottleneck trade-offs. Dell and QLogic offer dynamic bandwidth allocation and multiprotocol support in a next-generation converged network adapter.

Streamlining operations and minimizing the costs associated with running a data center are top-of-mind concerns for many organizations—particularly when a lagging economy leads to staffing limitations and budget constraints. Organizations can rise above such challenges and remain competitive by maintaining high quality of service (QoS) through innovative strategies designed to maximize IT resources and run applications efficiently.

Virtualization and convergence represent two key approaches that help organizations reduce costs by optimizing data center flexibility and efficiency. Virtualization helps reduce network infrastructure costs by facilitating the placement of application workloads on virtual machines rather than on dedicated physical servers. Recent advances in server platforms such as Dell™ PowerEdge™ 12th-generation servers enable IT administrators to move rising numbers

of virtual machines to physical servers. However, while enhancements in processor technology are designed to manage a growing number of virtual machine workloads, server I/O and potential bottlenecks pose significant QoS challenges for data center operations.

Implementation of 10 Gigabit Ethernet (10GbE) networking helps IT organizations avoid potential bottlenecks. But data center networks are being challenged by the need to provision bandwidth to specific virtual machine workloads and dynamically adjust bandwidth allocation to help ensure QoS for applications. In addition, convergence of traditional storage area network (SAN) and LAN protocols across 10GbE networks helps reduce network infrastructure costs. At the same time, convergence may lead to rising bandwidth requirements at the I/O level because newly converged networks need to support multiple protocols. To address these complexities, network administrators can deploy the QLogic QLE8262 10 Gbps Ethernet-to-PCI Express (PCIe) converged network adapter (CNA).

Provisioning and adjusting bandwidth allocation

Virtualization and network convergence contribute to cost-effective network consolidation, and 10GbE throughput helps accommodate increased I/O. However, these enhancements present two major challenges for data center networks. Administrators need to provision bandwidth to specific virtual machine workloads, and they need to dynamically adjust that bandwidth allocation for optimal network utilization to help ensure application QoS.

Dell and QLogic offer Switch Independent Partitioning that provides a method for provisioning bandwidth across virtual I/O data paths when deployed in selected Dell PowerEdge 12th-generation servers. QLogic enhances this capability through the use of multiprotocol support—including Fibre Channel over Ethernet (FCoE), Internet SCSI (iSCSI), and TCP/IP—across a single 10GbE adapter.¹ Once the virtual I/O paths have been established, administrators need a

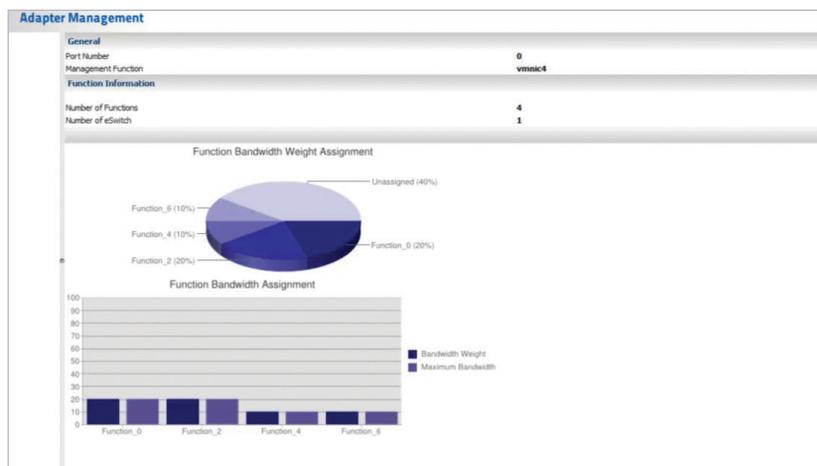


Figure 1. Using the vCenter plug-in to map virtual functions to virtual machines dedicated to specific applications

method to manage the bandwidth allocations to provide application workload QoS.

Managing bandwidth for quality of service

Administrators can use the VMware® vCenter™ plug-in from QLogic to manage bandwidth and QoS for the QLogic QLE8262 Ethernet-to-PCIe CNA. The plug-in includes settings to adjust bandwidth allocations in real time without requiring a server reboot to enable QoS for specific application workloads.

For example, in a scenario that includes four virtual machines dedicated to specific applications, QoS settings can address performance issues for one virtual machine running a Web server. In this scenario, a VMware ESX® 5 hypervisor host has four application servers installed: one running Microsoft® Exchange messaging software, one running Microsoft SharePoint® Server collaboration and documentation software, one running file server software, and one running the Web server software.

All four virtualized servers share a 10GbE pipe for client application access and physical port 0 on the QLE8262 Ethernet-to-PCIe CNA. In this case, the administrator can utilize Switch Independent Partitioning to map each server to its corresponding virtual function. For example, the Exchange virtual machine maps to Function_0,

¹ For more information, see "Enhancing I/O in a virtualized environment through NIC partitioning," by Michael Coon, Pankaj Gupta, and Rahul Deshmukh, in *Dell Power Solutions*, 2011 Issue 04, qrs.ly/l2g2cjs.



the SharePoint Server virtual machine maps to Function_2, the file server virtual machine maps to Function_4, and the Web server virtual machine maps to Function_6 (see Figure 1). Here, the clients connect to the applications using a Data Center Bridging (DCB)-capable switch and 10 Gbps connectivity. DCB is a set of standards-based extensions for Ethernet.

In this scenario, an administrator can allocate 20 percent of the available 10GbE bandwidth to Function_0 for the Exchange virtual machine and 20 percent to Function_2 for the SharePoint Server virtual machine. Similarly, a 10 percent allocation can go to Function_4 for the file server virtual machine and 10 percent can go to Function_6 for the Web server virtual machine. The remaining 40 percent of the bandwidth in this example remains unassigned and can be used for future growth and expansion.

If the administrator in this use case were to use the Microsoft Windows® OS-based Task Manager utility, which monitors network resources, it may show that Exchange is using approximately 16 percent of the bandwidth, which is within its 20 percent allotment. It may also show SharePoint Server using 16–18 percent of the bandwidth—within its 20 percent allotment—and the file server using 4 percent, which is also within its 10 percent allotment of bandwidth.

Because the Web presence for this example organization has experienced recent growth, the administrator may also note that the Web server virtual machine is running at 10 percent utilization. Because 10 percent is the allotment maximum, end users may notice a slowdown in Web access. For that reason, the Web server application in this example would warrant additional network bandwidth allotment to maintain QoS given the rise in access demand.

Taking this example to the next step, an administrator can accommodate the bandwidth uptick by changing the minimum and maximum allocation settings from 10 percent to 30 percent for the Web server

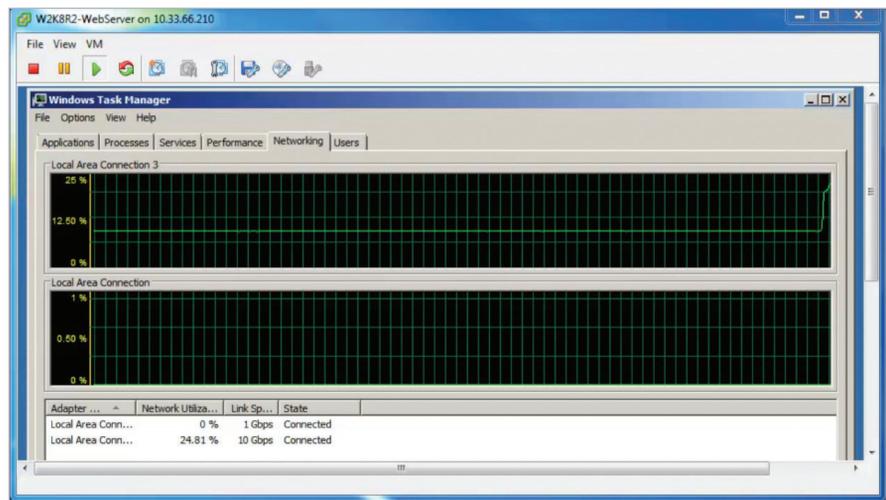


Figure 2. Utilizing bandwidth within a specified allocation following a dynamic increase adjustment

virtual machine mapped to Function_6. This dynamic change in bandwidth allocation helps ensure that the Web server has adequate bandwidth to handle current and rising levels of Web traffic. When the administrator saves these adjusted settings, the changes are implemented immediately and do not require a system reboot.

As a result of the bandwidth management described in this example, the administrator monitoring utilization of the Web server network may find that the Web server is using just under 24 percent of the bandwidth, which is well within its revised 30 percent allocation (see Figure 2). As this use case demonstrates, using the VMware vCenter plug-in to change a simple QoS parameter helps the administrator ensure that a mission-critical Web application has enough network bandwidth to continue providing optimal service to end users.

Allocating bandwidth for dynamic network I/O

Organizations looking to benefit from advances in virtualization and network convergence often need to address I/O bottleneck trade-offs for optimal throughput in their data center environment. Dell and QLogic have partnered to help enhance bandwidth

utilization and application QoS through Switch Independent Partitioning that provisions bandwidth across virtual I/O data paths. By deploying the QLogic QLE8262 Ethernet-to-PCIe CNA in selected Dell PowerEdge 12th-generation servers, administrators can take advantage of multiprotocol support and bandwidth allocation and adjustment to enhance application QoS. In addition, the VMware vCenter plug-in from QLogic lets administrators easily manage bandwidth allocation and make necessary adjustments to maintain high levels of QoS for mission-critical applications. **PS**

Authors

Kyle Roblyer is a senior global marketing manager at QLogic. He has more than 10 years of experience working with Dell to drive adoption of leading-edge networking solutions.

Rahul Deshmukh is a senior technical marketing manager at Dell specializing in server peripheral devices.

Learn more

Dell and QLogic:
qlogic.com/go/dell

Dell and QLogic solutions:
qrs.ly/762gcjl