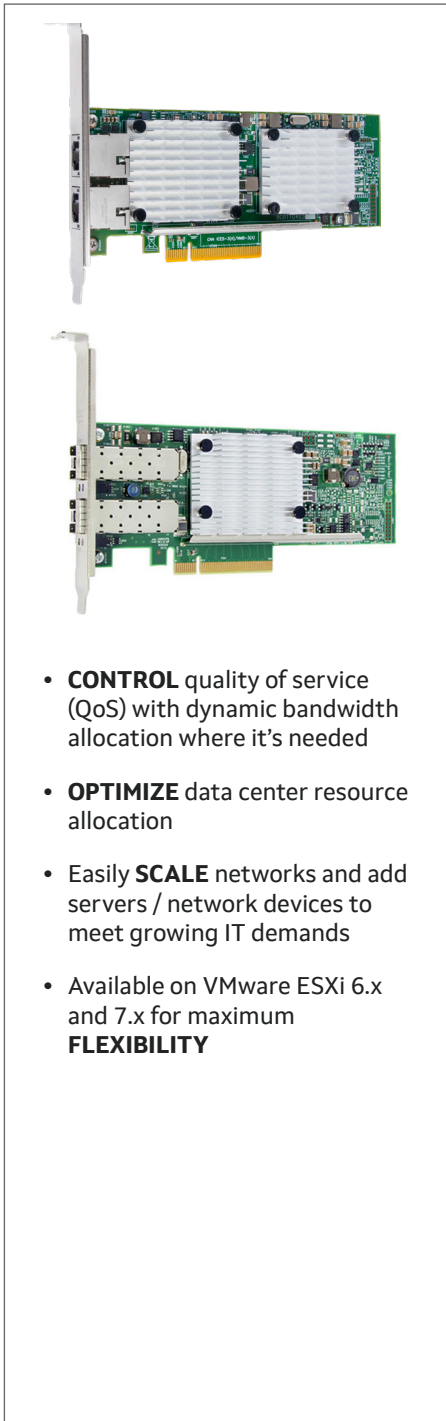


Guarantee Application Performance in VMware with Marvell NIC Partitioning

Practical Use Cases with Marvell FastLinQ Ethernet Adapters



- **CONTROL** quality of service (QoS) with dynamic bandwidth allocation where it's needed
- **OPTIMIZE** data center resource allocation
- Easily **SCALE** networks and add servers / network devices to meet growing IT demands
- Available on VMware ESXi 6.x and 7.x for maximum **FLEXIBILITY**

Introduction

Data centers continue to demand a high degree of efficiency with hyper-scale virtualized server deployments. Applications, rich media, software-defined networking, and cloud services drive enterprises to migrate to 10/25Gb Ethernet to improve the available bandwidth and scalability for virtualized hosts and services. However, the migration also introduces efficiency challenges when using 10/25Gb bandwidth.

Designed for data center I/O efficiency and flexibility, Marvell Switch Independent NIC Partitioning (NPAR) technology addresses the challenges of efficient bandwidth utilization by maximizing unused bandwidth with flexible I/O NPAR. This simplifies end-to-end networking by dividing a network controller port into as many as four partitions, enabling dynamic allocation of bandwidth as needed while reducing the total cost of ownership. NPAR is switch agnostic, supported across multiple operating systems, and provides seamless integration with easy-to-use management options.

This paper provides specific examples of using NPAR in VMware® virtual environments. This includes two main use cases: the separation of traditional production VMware environments and the co-mingling of production and development/testing environments.

First Use Case – Separation Of Traditional VMware Traffic

A typical VMware environment includes management traffic, vMotion traffic, and Virtual Machine (VM) traffic. It might also support fault tolerant and software iSCSI traffic. Focusing on just the first three, there are two ways to achieve the separation of traffic: the use of VLANs and/or separate physical network connection and vSwitches. While VLANs provide separation and security, QoS at the switch level requires switch and licensing to support. Physical separation gives the best flexibility since it can be tailored to the needs of the environment like using 1GbE for management and 10GbE for vMotion and VM traffic. However, this requires using a larger number of physical adapters and cabling.

NPAR can deliver the best of both worlds by using two dual-port Marvell FastLinQ adapters and partitioning their ports based on different traffic demand requirements. This not only provides the ability to have separate physical functions (PFs) that can be assigned to different vSwitches, but with NPAR you can set the QoS for each function without having to modify the network configuration.

Guarantee Application Performance in VMware with Marvell NIC Partitioning Solution Sheet

The easiest configuration to achieve this level of separation is to partition port 0 on each card into two partitions. These partitions are used for the management and vMotion traffic. Port 1 on each adapter is used exclusively for VM traffic. This basic configuration provides redundancy of the physical adapters by using one port from each to connect the vSwitches. The table below illustrates the settings for these ports.

Port/Partition	Use	QoS Setting	vSwitch Assignment
Port 0/Partition 1	Management network traffic	10%	Management vSwitch
Port 0/Partition 2	vMotion network traffic	90%	vMotion vSwitch
Port 1	VM network traffic	100%	VM vSwitch

In this scenario, 1Gb is dedicated bandwidth for management, 9Gb for vMotion, and 10Gb for VM traffic per NIC port.

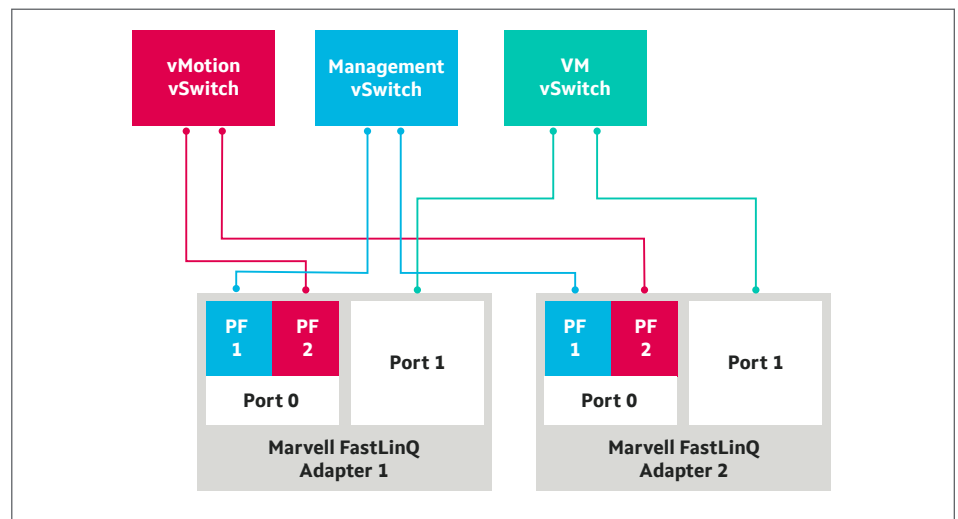


Figure 1. Separation of Traditional VMware Traffic

Flexible And Optimized Bandwidth Allocation

Each partition is configured with relative bandwidth weight and maximum bandwidth. The maximum bandwidth limits the bandwidth given to a partition, even if there is available bandwidth that no other partition uses. The relative weight determines the relative bandwidth a partition has with regard to the other partitions on the same port when they compete for the same available bandwidth. A partition obtains its relative weight from the available bandwidth, but never more than its configured maximum bandwidth. If a partition does not use its configured weight, its bandwidth becomes available for other partitions to use.

Another popular option is to use the lower bandwidth on-board 1GbE NIC ports for management traffic. Port 0 on each of the two adapters could be divided into two partitions. In this use case, partition 1 would be used only for Fault Tolerant Logging and partition 2 would be for vMotion network traffic. This leaves port 1 on each adapter to be 100% dedicated to VM network traffic.

Guarantee Application Performance in VMware with Marvell NIC Partitioning Solution Sheet

Port/Partition	Use	QoS Setting	vSwitch Assignment
1GbE LAN On Motherboard	Management traffic	NA	Management vSwitch
Port 0/ Partition 1	Fault Tolerant Logging	50%	FT vSwitch
Port 0/ Partition 2	vMotion network traffic	50%	vMotion vSwitch
Port 1	VM network traffic	100%	VM vSwitch

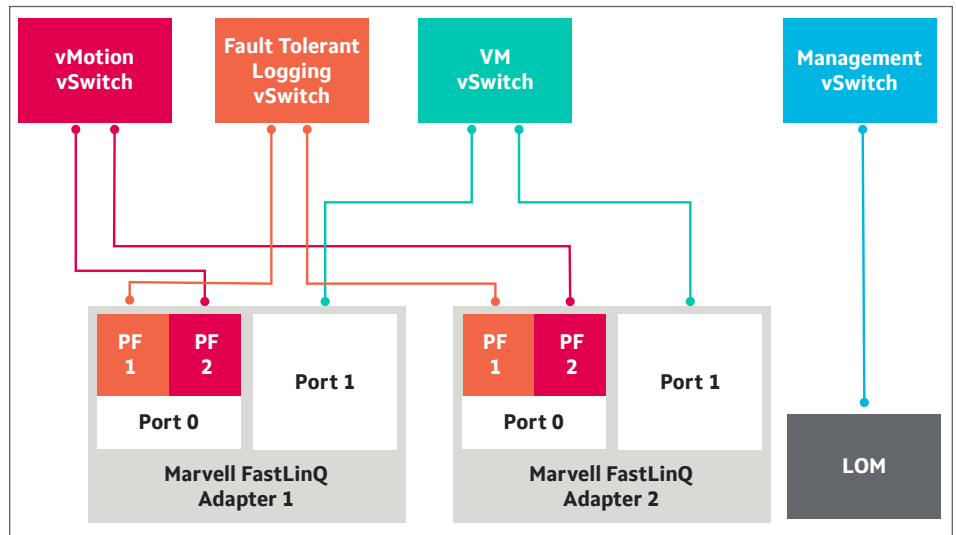


Figure 2. Separation of Traditional VMware Traffic with LAN on Motherboard (LOM)

Second Use Case – Production And Development/Testing In The Same Environment

Expanding on the first use case, it is possible to run both production and development/testing environments in the same VMware environment. However, the development/testing environment can introduce performance or latency issues in the production environment. Administrators can easily overcome this by partitioning the second port on the adapter to limit the development/testing traffic and provide a guaranteed bandwidth for the production traffic.

Using a dual-port adapter in this scenario, the administrator can divide one 10GbE port into one partition with 1Gb dedicated bandwidth for management and 9Gb for vMotion traffic; and the second port into a partition with 6Gb guaranteed for production VM traffic. and the other with 4Gb for development/testing VM traffic.

Port/Partition	Use	QoS Setting	vSwitch Assignment
Port 0/Partition 1	Management network traffic	10%	Management vSwitch
Port 0/Partition 2	vMotion network traffic	90%	vMotion vSwitch
Port 1/Partition 1	Production VM network traffic	60%	Production VM vSwitch
Port 1/Partition 2	Test/Dev VM network traffic	40%	Test/Dev VM vSwitch

SUMMARY

10/25 GbE networks can deliver the flexibility that enhances the efficiency of your servers and storage. Switch Independent NIC Partitioning makes the migration to today's powerful 10GbE networks an easy and compelling option for many organizations. With optimized I/O and server resource utilization, performance and efficiency is maximized while reducing the total cost of ownership (TCO).

- **Reduced Network Sprawl:** With a two-port GbE device, a data center can usually deploy 10/25 NICs, two host bus adapters, and many cables and switch ports to meet bandwidth needs. With the 10/25GbE Switch Independent NIC Partitioning approach, the same capacity and flexibility can be provided with a 10/25GbE Converged Network Adapter—consolidating the infrastructure onto much less equipment while using less facility space than when deploying GbE devices.
- **Maximized Network Scalability:** Having a reduced number of network devices and cables allows IT organizations and data centers to easily scale their networks and add servers and network devices to meet growing IT demands.
- **Simplified Administration:** Switch Independent NIC Partitioning can also save time and labor by helping to simplify the management tasks for IT administrators and enhance their self-sufficiency. They can add or replace network cards or move workloads from one partition to another within minutes.
- **Optimized Resource Allocation:** Switch Independent NIC Partitioning can play an important role in helping to optimize bandwidth allocation and the effective utilization of that allocation in both virtualized and non-virtualized environments. Virtualization can reduce the number of physical servers and increase I/O demand for each server. Using 10/25GbE connectivity with Switch Independent NIC Partitioning addresses this problem by providing eight functions on a single integrated network adapter built into the server, so the server can efficiently handle current bandwidth requirements and still allow additional devices to be plugged in to meet future bandwidth needs. Switch Independent NIC Partitioning also allows system administrators to fine-tune the amount of bandwidth for each adapter port, as required by the application, and to set the bandwidth for each physical function. Switch Independent NIC Partitioning helps optimize enterprise data centers by providing these benefits for 10/25GbE networks.



To deliver the data infrastructure technology that connects the world, we're building solutions on the most powerful foundation: our partnerships with our customers. Trusted by the world's leading technology companies for 25 years, we move, store, process and secure the world's data with semiconductor solutions designed for our customers' current needs and future ambitions. Through a process of deep collaboration and transparency, we're ultimately changing the way tomorrow's enterprise, cloud, automotive, and carrier architectures transform—for the better.

Copyright © 2020 Marvell. All rights reserved. Marvell and the Marvell logo are trademarks of Marvell or its affiliates. Please visit www.marvell.com for a complete list of Marvell trademarks. Other names and brands may be claimed as the property of others.