


NIC Partitioning and SR-IOV

Cavium is Driving Next-Generation Server I/O Virtualization with NIC Partitioning



NIC Partitioning (NPAR) together with SR-IOV provides a comprehensive solution for the heterogeneous enterprise data center.

INDUSTRY CHALLENGES

Virtualization of networks, storage, servers, and applications continues to grow in the data center. While virtualization provides many benefits, virtual deployment workloads increase the demands placed on the network I/O infrastructure, which can significantly impact system and application performance.

Architectural improvements occur on the interface between the CPUs and the network interface hardware. Virtualization of I/O itself using Single Root I/O Virtualization (SR-IOV) for individual servers and Multi Root I/O Virtualization (MR-IOV) for blade servers (which share I/O resources) is an emerging technology standard that promises to improve the scalability, flexibility, throughput, and latency performance of networking operations.

NIC Partitioning (NPAR) is similar to SR-IOV in that it logically divides a physical adapter port into multiple virtual ports. SR-IOV lacks broad support today by the majority of popular operating systems (OSs). To deploy SR-IOV today, an organization needs to ensure a minimum level of infrastructure (server hardware and OS) support for SR-IOV. In contrast, NPAR is available today with the Cavium™ Converged Network Adapters and supported with all major OSs, including Windows®, VMware®, and Linux®.

NIC PARTITIONING – HARDWARE-BASED QoS

Cavium is driving next-generation server I/O virtualization with NPAR, which is based on the QLogic® VMflex™ technology. NPAR is a method of dividing a single physical 10GbE Ethernet port into multiple PCI physical functions or partitions with flexible bandwidth capacity allocation. This approach enables the application of QoS to the virtual ports, improving I/O performance while maintaining a low total cost of ownership (TCO).

Cavium's implementation maps four PCI functions to each physical port on a dual-port 10GbE device. NPAR presents the eight PCI functions per device using standard PCI configuration space. Each function or partition is assigned a unique MAC address. Each partition can support concurrent networking or storage protocols, enabling flexible bandwidth provisioning to applications.

Unlike competitive solutions, these partitions can be assigned to run both networking and storage protocols with an implementation that is agnostic to the external Ethernet switch. For certain configurations, this approach also enables the switching of VM-to-VM traffic via the embedded switch located in the adapter. This results in lower CPU utilization while preserving I/O performance, providing sufficient capacity to run additional applications. NPAR provides a non-captive (switch agnostic) solution to customers, thereby eliminating vendor lock-in.

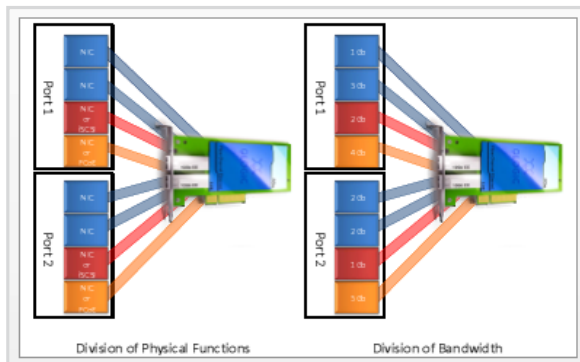


Figure 1. NPAR

PCI-SIG SR-IOV STANDARD

PCI special interest group (PCI-SIG®) has completed a suite of I/O virtualization specifications that in conjunction with system virtualization technologies allow multiple OSs running simultaneously within a single computer to natively share PCI Express® devices.

I/O virtualization is an emerging category of technologies aimed at resolving the performance, provisioning, and management problems of server I/O in highly virtualized and clustered data center environments. In the evolution of the virtualized data center, the issues of CPU and storage virtualization are actively being addressed, but are raising problems with the load, provisioning, and management of storage and network infrastructure. The technology evolution started with virtualized memory. Next came virtual machines (VMs), then virtual storage, and finally I/O virtualization (IOV), where the I/O path from the server to the peripheral device is now virtualized.

The SR-IOV specification designates how a hardware device can expose multiple “light-weight” Virtual Functions (VFs) for use by VMs. VFs are focused primarily on data movement and are associated with a Physical Function (PF). The PF is equivalent to the regular PCIe® physical function device’s fully featured PCI functions. The PF is responsible for arbitration relating to policy decisions (such as link speed or MAC addresses in use by VMs in the case of networking) and for I/O from the parent partition itself. The PF is strictly owned by the resident kernel; therefore changes to a PF’s personality require a system reboot. Because the VFs use hardware resources, there are constraints on the number of VFs that are available on any particular device. There may be several VFs associated with a single PF. The VFs provide a function of the type defined by the PF with which it is associated. A NIC PF can have several NIC VFs, for example. Or an FCoE PF can have several FCoE VFs. A VF must be associated with a PF. The VFs are under the control of the OS. The OS then is able to make or break the association of a VF to a VM without having to reboot. Although the SR-IOV standard applies to networking and storage I/O, the initial implementations are for networking I/O only.

With NPAR, each physical port is partitioned into multiple physical functions on the PCIe bus. However, in the case of SR-IOV, the physical ports are further partitioned into multiple virtual functions. This difference allows NPAR to be deployed in both bare metal (non-virtualized) OSs and virtualized OSs. In contrast, SR-IOV is primarily targeted towards virtualized platforms.

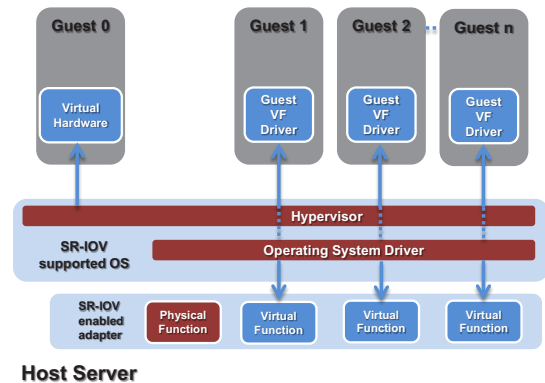


Figure 2. SR-IOV

KEY BENEFITS

Here are some of the ways NPAR offers value to an organization:

Reduced Capital and Operational Expenditures

- Enables hardware consolidation of multiple 1GbE ports into a single 10GbE port
- Allows the creation of multiple partitions for flexible bandwidth provisioning on 10GbE adapters
- Provides lower adapter, cabling, switch port, and management costs

Eliminates Dependency on a Captive Switch

- Removes dependency on an external Ethernet switch to regulate and manage bandwidth
- Offers customers the freedom to choose an Ethernet switch of their choice

Simplified Deployment and Management

- Provides concurrent support for Ethernet, FCoE, and iSCSI protocols, which appear as discrete functions similar to native OS, minimizing deployment disruptions
- Eliminates the need for OS changes to implement flexible bandwidth provisioning

Improved Ability to Scale the Business

- Reduces I/O emulation overhead
- Assigns I/O hardware directly to VMs
- Allows the dedication of bandwidth (QoS) for VMs and associated applications
- Conserves PCIe slots in all server environments
- Enables high I/O performance in virtualized environments

Advantages of NPAR

- Available today

SUMMARY

SR-IOV is a planned standard soon to be recognized by most OS and hypervisor vendors. It enables virtualization of I/O data paths at the server level. Cavium offers NPAR as a standard way of addressing the customer needs of today for the network partitioning function while the SR-IOV technology is under deployment. Cavium NPAR is an OS-, platform-, and switch-agnostic solution available today. NPAR together with SRIOV provides a comprehensive solution for the heterogeneous enterprise data center.

ABOUT CAVIUM

Cavium, Inc. (NASDAQ: CAVM), offers a broad portfolio of infrastructure solutions for compute, security, storage, switching, connectivity and baseband processing. Cavium's highly integrated multi-core SoC products deliver software compatible solutions across low to high performance points enabling secure and intelligent functionality in Enterprise, Data Center and Service Provider Equipment. Cavium processors and solutions are supported by an extensive ecosystem of operating systems, tools, application stacks, hardware reference designs and other products. Cavium is headquartered in San Jose, CA with design centers in California, Massachusetts, India, Israel, China and Taiwan.



Follow us:      

Corporate Headquarters Cavium, Inc. 2315 N. First Street San Jose, CA 95131 408-943-7100

International Offices UK | Ireland | Germany | France | India | Japan | China | Hong Kong | Singapore | Taiwan | Israel

Copyright © 2012 - 2017 Cavium, Inc. All rights reserved worldwide. QLogic LLC (formerly QLogic Corporation) is a wholly owned subsidiary of Cavium, Inc. Cavium, QLogic, QConvergeConsole, and VMflex are registered trademarks or trademarks of Cavium Inc., registered in the United States and other countries. All other brand and product names are registered trademarks or trademarks of their respective owners.

This document is provided for informational purposes only and may contain errors. Cavium reserves the right, without notice, to make changes to this document or in product design or specifications. Cavium disclaims any warranty of any kind, expressed or implied, and does not guarantee that any results or performance described in the document will be achieved by you. All statements regarding Cavium's future direction and intent are subject to change or withdrawal without notice and represent goals and objectives only.