

NIC Partitioning and Data Center Bridging

Cavium Delivers iSCSI over Data Center Bridging with Type-Length-Value (TLV) Support



Data Center Bridging (DCB) is a series of enhancements to the Ethernet protocol that enable the use of traditional Ethernet-based LAN switching to support Fibre Channel SAN data center traffic.

INTRODUCTION

Data Center Bridging (DCB) is a series of enhancements to the Ethernet protocol that enable the use of traditional Ethernet-based LAN switching to support Fibre Channel SAN data center traffic.

This technology brief examines the use of Cavium™ Converged Network Adapters with NIC Partitioning (NPAR) to enhance the user flexibility of managing virtual data paths at the I/O level in a DCB environment with iSCSI TLV. Converged networking requires enhanced Ethernet capability with quality of service (QoS) support to deliver the reliable solutions that enterprise data center deployments require.

Guaranteed granular bandwidth allocation for storage and networking traffic of a 10Gb Ethernet is user configurable when QoS elements are supported. QoS settings allow the allocation of bandwidth between traffic classes and assign priority attributes, providing efficient I/O utilization and management of networked traffic and eliminating the need for discrete 1Gb Ethernet ports—enhancing the value of network virtualization.

NPAR: 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). The solution 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 and storage protocols, enabling flexible bandwidth provisioning to applications. Unlike competitive solutions, these partitions can be assigned to run both networking or storage protocols, with an implementation that is agnostic to the external Ethernet switch. For certain configurations, the unique eSwitch implementation enables the switching of VM-to-VM traffic via the embedded switch located in the adapter, thereby reducing the traffic and latency of flow through an external switch port. This results in lower CPU utilization while preserving I/O performance, providing sufficient capacity to run additional applications.

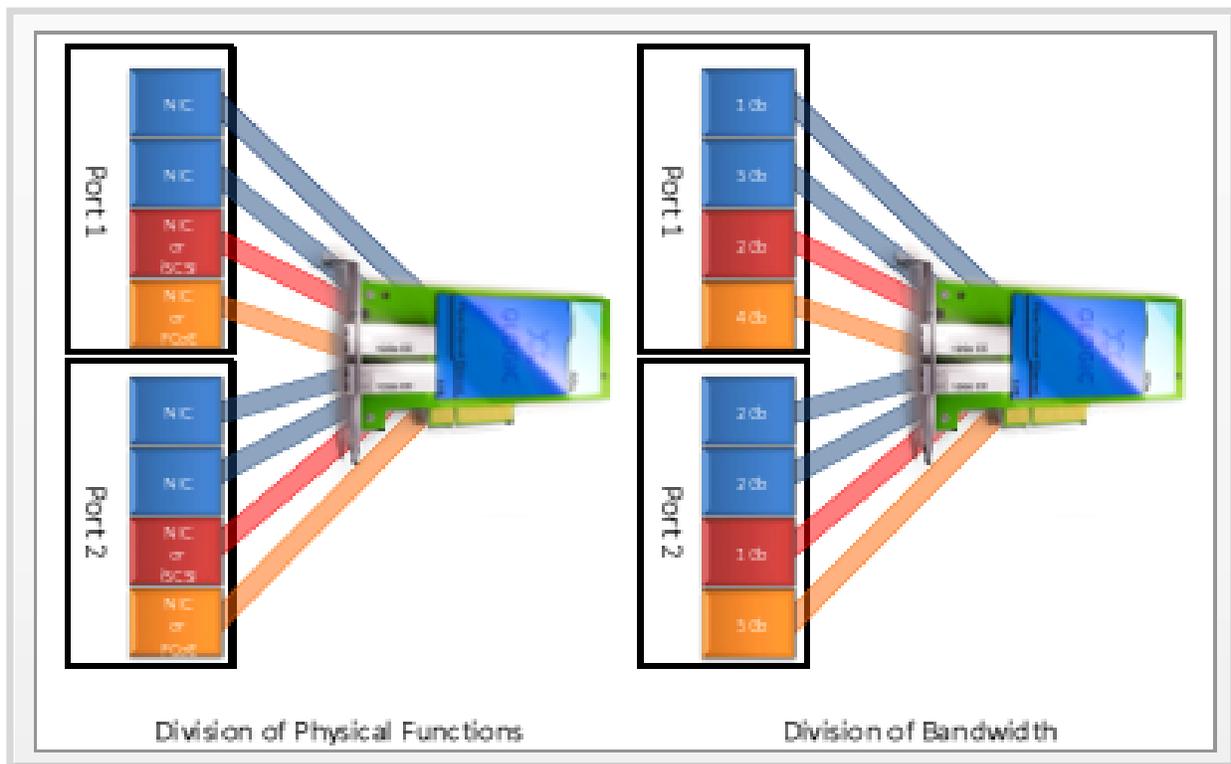


Figure 1. NPAR

NPAR coupled with the unique eSwitch implementation provides a non-captive (switch and OS agnostic) solution to customers, ensuring the highest levels of virtualization with assured interoperability in heterogeneous data centers. (See Figure 1.)

DATA CENTER BRIDGING

Lossless Ethernet is specifically designed to meet the needs of storage applications by managing the flow of traffic onto the network to avoid congestion and the risk of packet loss for the storage traffic. DCB encompasses four new capabilities that were added to the Ethernet 802.1Q specifications to support new requirements to allow storage traffic to be transported over Ethernet: Priority-based Flow Control (PFC: 802.1Qbb) for controlling the amount of traffic that is allowed onto the network based upon the defined traffic class, Enhanced Transmission Selection (ETS: 802.1Qaz) for classifying traffic and requesting lossless services, Congestion Notification (CN: 802.1Qau) for monitoring the congestion in the network and triggering flow control events between devices, and an exchange protocol called DCB exchange (DCBx: 802.1Qaz) for devices that support DCB to communicate capabilities and configuration information to each other and manage the network.

This protocol leverages functionality provided by IEEE® 802.1AB (LLDP) and is included in the 802.1az standard. DCB also leverages existing capabilities of standard Ethernet devices to dedicate specific amounts of bandwidth to specific traffic classes and enforce class-based priorit access to the network. DCBx is used to exchange configuration information with the directly connected peer. DCBx is expected to operate only over a point-to-point link and if multiple peers are discovered, the peers' TLVs should be ignored until the multiple peers' condition is resolved.

Cavium was the first to introduce the Converged Network Adapter class of converged networking products to the market and maintains a leadership position in that segment. The Cavium Converged Network Adapters support concurrent LAN (TCP/IP) and SAN (FCoE and iSCSI) traffic over a shared 10Gb Ethernet link and are ideal for beginning the evolution to converged networking at the SAN edge. The adapters deliver industry-leading LAN networking performance, and they are the only adapter offering in the market with full hardware offload for FCoE and iSCSI protocol processing and an extensive set of features that enhance system virtualization.

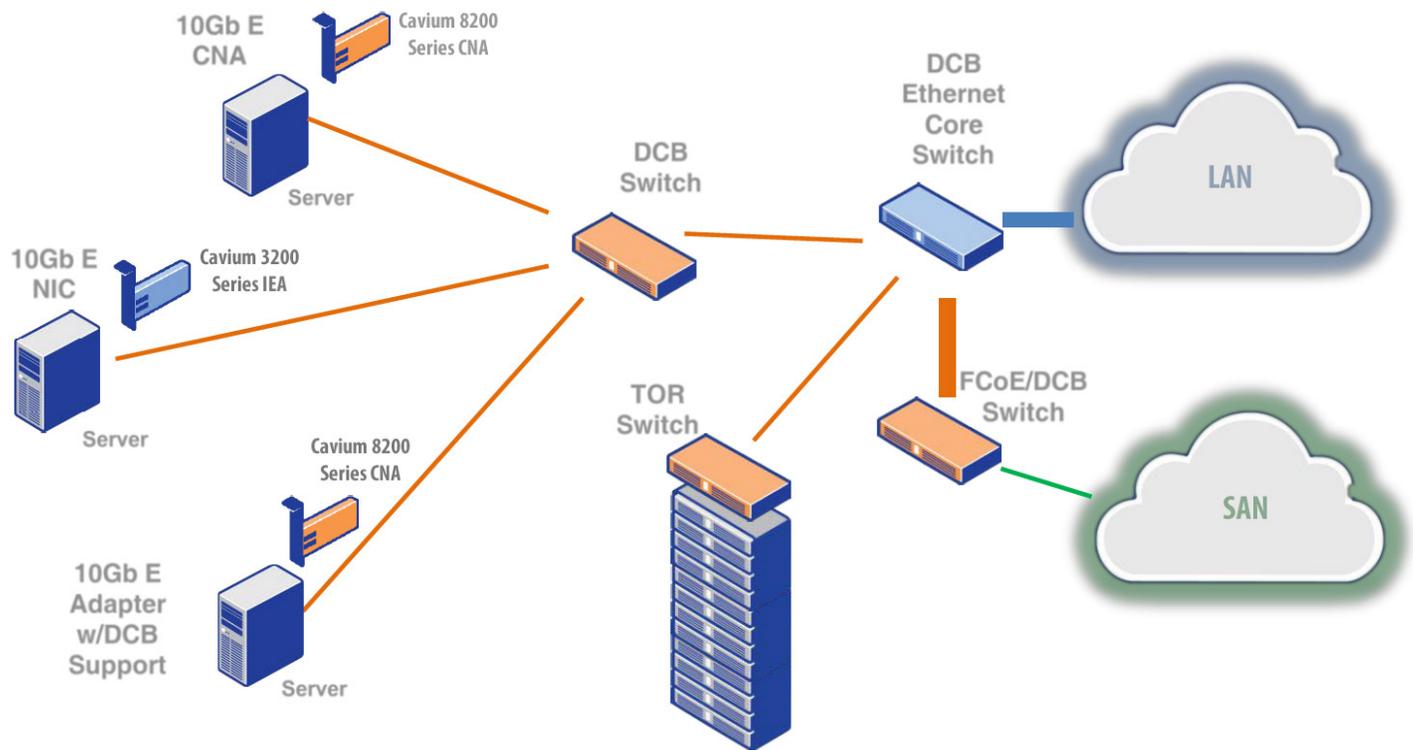


Figure 2. Data Center Bridging

In addition, these adapters leverage existing QLogic Fibre Channel and iSCSI drivers that have gone through more than 10,000 man-months of reliability testing and are field hardened with deployments of more than 13 million ports in enterprise data centers, making the adapters the most dependable and field-hardened Converged Network Adapters on the market. (See Figure 2.)

ISCSI OVER DCB

Supporting DCB for iSCSI allows the Cavium 10Gb iSCSI to participate in Enhanced Transmission Selection (ETS), which allocates to iSCSI network bandwidth, and Priority Flow Control (PFC), which prevents dropped frames on a per-protocol basis. This implementation allows iSCSI and other traffic (NIC or FCoE) to converge on the same cable, each having an allocated share of the bandwidth and a PFC priority with which it can be independently throttled without affecting the flow of the other protocols' traffic.

DCB AND NPAR

ETS enables a converged network switch to allocate bandwidth to the converged protocols (NIC, FCoE, and iSCSI) configured on an attached port. Traffic transmitted from the Converged Network Adapter to the switch is examined by the switch that passes the traffic to the fabric proportionally to the protocols' allocations established by DCBx. If a Converged Network Adapter port supports NPAR and is configured for multiple NIC functions, ETS makes no distinction between the traffic flows from the different NIC functions.

Cavium provides NPAR QoS to divide the NIC bandwidth among the configured NIC functions. The DCBx specifies the application protocol TLV that allows the upper layer protocols and related priority information to be advertised and exchanged over the DCB link. The iSCSI TLV application protocol parameters are used to set the priority flow control and any other required parameters. A traffic agreement is set between the two nodes, and all traffic subsequent to the TLV agreement follows the settings established in the exchange.

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 virtual machines (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

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, 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.