

# Cavium FastLinQ BCM57840S

10Gbps Quad-Port iSCSI, FCoE, and TOE PCI-SIG SR-IOV x8 PCI Express 3.0 Converged Controller



- Low-power, single-chip solution for four ports of 10GBASE-KR compliant backplane 1G/10G Ethernet
- Low-power, single-chip solution for four ports of SFP+ optical Converged Network Adapter
- iSCSI v1.0 Host Bus Adapter
- FCoE Host Bus Adapter applications
- Applications with Energy Efficient Ethernet (EEE) for power savings with 10GBASE-T
- Virtualization environments
- 10GBASE-KR mezzanine card for server blades
- 10GBASE-T copper NIC with external 100M/1G/10G copper PHY (BCM84833, for example)

## OVERVIEW

The Cavium™ FastLinQ® BCM57840S Converged Controller is a sixth-generation converged controller designed for high-volume, converged LAN on motherboard (LOM) and Converged Network Adapter applications. The controller enables PCI-SIG single root input/output virtualization (SR-IOV), iSCSI, Fibre Channel over Ethernet (FCoE), data center bridging (DCB), and on-chip TCP/IP offload engine (TOE). The converged controller supports PCI Express® (PCIe®) 3.0, along with embedded virtual bridging and other switching technologies for high-performance DMA and virtual machine (VM)-to-VM switching.

The FastLinQ BCM57840S Controller includes quad-channel 10GBASE-KR and SFF-8431 for SFP+ 10Gb and SFP 1Gb interfaces. The FastLinQ BCM57840S Controller integrates four IEEE 802.3™-compliant MACs and supports the network controller-sideband interface (NC-SI). Host-BMC communication is also supported on top of NC-SI to permit high-speed communication between the local host and the baseboard management controller (BMC) or management controller (MC). The feature-complete converged controller requires only 0.82 square inches of printed circuit board (PCB) space and enables 10G speeds at low per-port power.

For more effective use of 10GbE bandwidth, the FastLinQ BCM57840S Controller offers Cavium switch-independent NIC partitioning (NPAR), which enables the segmentation of a single port into two virtual ports (in 4 × 10G mode) or four virtual ports (in 2 × 10G or 2 × 20G modes) with

flexible bandwidth allocation to each partition. The segmentation allows IT organizations to improve resource usage while lowering infrastructure and operational costs.

The FastLinQ BCM57840S Controller enables convergence of all possible network communications in a server, such as data network (LAN), FCoE storage network or block (for example, iSCSI). The FastLinQ BCM57840S Controller can simultaneously support all offload traffic types on each of the ports, including simultaneous iSCSI and FCoE. Offload results in superior storage and networking performance, as well as low CPU usage, which results in significant system-level power savings.

The FastLinQ BCM57840S Controller is designed for PCIe 3.0 and is also compatible with *PCI Express Base Specification*, revisions 2.0 and 1.1. PCI Express supports MSI and MSI-X capabilities. A separate PCI function is supported for each of the ports.

## FEATURES

### Media Interfaces

- Integrated quad 10Gbps MAC with offload and quad 10GBASE-KR/SFF-8431 (SFP+)
- Single 25.00MHz clock crystal for quad-port 10Gbps operation
- 2×10G, 4×10G, 2×20G modes

**Host Interfaces**

- PCIe x8 3.0, 8GT/s and x8 2.0, 5GT/s compliant
- PCIe lanes x1, x4, and x8
- No external dynamic random-access memory (DRAM) required; flow-through architecture
- PCIe CLKREQ support
- SR-IOV
- Comprehensive IPv4 and IPv6 stateless offloads
- Broad OS and hypervisor support
- Receive side scaling (RSS) and transmit side scaling (TSS)
- Support for jumbo frames up to 9,600 bytes
- Network teaming, failover and load balancing
- MSI and MSI-X support
- Switch-independent NPAR
- Generic routing encapsulation (NVGRE) packet task offloads
- Virtual extensible LAN (VXLAN) packet task offloads
- Generic network virtualization encapsulation packet task offloads
- Generic routing encapsulation (GRE) packet task offloads

**Network Interfaces**

- Quad-port 10GBASE-KR/SFF-8431 (SFP+) interfaces for 1Gbps and 10Gbps operation
- IEEE 802.3ap Clause 73-compliant backplane operation
- IEEE 802.3xx Clause 37-compliant auto-negotiation for 1Gbps

**TCP/IP Offload Engine (TOE)**

- Microsoft® TCP chimney compliant

**iSCSI Controller**

- Offloaded full Host Bus Adapter function iSCSI initiator
- iSCSI boot and iSCSI crash dump support

**FCoE**

- Receiver and transmitter CRC offload
- Offloaded full Host Bus Adapter function FCoE initiator
- FCoE boot from SAN
- Large, concurrent port logins and exchanges (4,096 each)
- N\_Port ID virtualization (NPIV)
- Virtual Fibre Channel (vFC) on Windows Server 2012 and later Hyper-V

**Robust Manageability**

- NC-SI
- Pre-execution environment (PXE) v2.1 remote boot
- Statistics gathering (SNMP management information base [MIB] II and Ethernet MIB) (IEEE 802.2x, Clause 30)
- Comprehensive diagnostic and configuration software suite

**DCB**

- Enhanced transmission selection (ETS) (IEEE 802.1Qaz)
- Quantized congestion notification (QCN)-capable (IEEE 802.1Qau)
- Priority-based flow control (PFC) (IEEE 802.1Qbb)
- Lossless iSCSI-Offload-TLV over DCB

**BENEFITS**

- SR-IOV 10Gbps and converged solution—Power and space optimized for blade server, rack, tower, and Converged Network Adapter applications
- Extremely low CPU usage for iSCSI, FCoE, and TCP/IP
  - Host CPU is free to run application code
  - Minimal load on memory subsystem with zero copy
- Accelerated IP-based file and block storage
  - Lower CPU usage for file-level storage protocols such as common Internet file system (CIFS), server message block (SMB) protocol, and NFS
  - Offloaded and accelerated iSCSI block storage with high I/O per second and low CPU usage
- Accelerated FCoE
  - Offloaded and accelerated FCoE for Fibre Channel block storage with high I/O per second and low CPU usage
- Performance-focused—Optimized for high throughput, low latency, and CPU usage
  - Adaptive interrupt coalescing
  - RSS reduces CPU usage on multi-CPU systems
  - MSI and MSI-X allows interrupt distribution in a multi-CPU system.
- Robust and highly manageable
  - NC-SI enables high bandwidth out-of-band system management over shared infrastructure.
  - Guaranteed delivery of management traffic
  - PXE v2.1, ACPI v2.0b, and WoL
  - Host-BMC communication for connectivity between local host and management controller (MC or BMC)
- Server class reliability, availability, and performance features
  - Link aggregation and load balancing (switch-dependent)
  - IEEE 802.3ad (link aggregation control protocol [LACP]), generic trunking (GEC/FEC) (switch- and NIC-independent)
- RoHS compliant

**PART NUMBER**

- B57840SB1KFSBR

## Host Bus Interface Specifications

### Bus Interface

- PCIe 3.0 x8 (x8 physical connector)

### Host Interrupts

- MSI-X supports independent queues

### I/O Virtualization

- SR-IOV (128 maximum virtual functions per device)
- Cavium switch-independent NPAR (8 physical function partitions)
- Network virtualization using generic routing encapsulation (NVGRE) packet task offloads
- Virtual extensible LAN (VXLAN) packet task offloads
- Generic network virtualization encapsulation (GENEVE) packet task offloads
- Generic routing encapsulation (GRE) packet task offloads

### Compliance

- PCI Express Base Specification, rev. 3.0
- PCI Bus Power Management Interface Specification, rev 1.2
- Advanced Configuration and Power Interface (ACPI), v2.0
- SMBus 2.0

## Ethernet Specifications

### Throughput

- 10Gbps full-duplex line rate per 10G port
- 20Gbps full-duplex line rate per 20G port

### Ethernet Frame

- Standard MTU sizes; jumbo frame up to 9,600 bytes

### Stateless Offload

- TCP segmentation offload (TSO)
- Large send offload (LSO)
- Large receive offload (LRO)
- Giant send offload (GSO)
- TCP and user datagram protocol (UDP) checksum offloads
- Hardware transparent packet aggregation (TPA)
- Receive segment coalescing (RSC)
- Interrupt coalescing
- RSS and TSS—Maximum of 16 queues per any (1GbE or 10GbE) physical function (PF) in single function (SF) and Cavium switch-independent partitioning modes
- VMware® NetQueue and Microsoft dynamic virtual machine queue (VMQ)

### Compliance

- IEEE 802.3ae (10Gb Ethernet)
- IEEE 802.1q (VLAN)
- IEEE 802.3ad (Link Aggregation)
- IEEE 802.3-20015 (Flow Control)
- IPv4 (RFC 791)
- IPv6 (RFC 2460)
- IEEE 802.1Qbb (Priority-Based Flow Control)
- IEEE 802.1Qaz (DCBX and Enhanced Transmission Selection)
- IEEE 802.1AS/1588-2002 PTPv1 (Hardware Precision Time Protocol)
- IEEE 1588-2008 PTPv2
- IEEE 802.3-2015 Clause 52 (10Gb Ethernet optical on SFP ports)
- SFF8431 Annex E (10Gb Direct Attach Copper on SFP ports)
- IEEE 802.3an-12 Clause 55 10GBASE-T (on 10GBASE-T ports)
- IEEE 802.3ab-2012 Clause 39 1000BASE-T (on BASE-T ports)
- IEEE 802.3-2012 Clause 25 100BASE-TX (on BASE-T ports)
- IEEE 802.3i 10BASE-TX (on 1000BASE-T ports)
- IEEE 802.3az (Energy Efficient Ethernet on BASE-T ports)
- SFF8431 (enhanced Small Form Factor Pluggable modules)

## Tools and Utilities

### Management Tools and Device Utilities

- QConvergeConsole® (QCC) integrated network management utility graphical user interface (GUI) for Linux® and Windows
- Cavium FastLinQ PowerKit cmdlets for Linux and Windows
- QCC Plug-in for vSphere® (GUI) and ESXCLI plug-in for VMware
- QLogic Control Suite (QCS) Command Line Interface (CLI) for Linux and Windows
- Pre-boot unified extensible firmware interface (UEFI) Device Configuration pages in system BIOS
- Cavium Comprehensive Configuration Management (CCM)
- Native OS management tools for networking

### Boot Support

- iSCSI remote boot
- FCoE boot from SAN
- PXE 2.0

### Operating System Support

- For the latest applicable operating system information, see [www.cavium.com](http://www.cavium.com), **Downloads**

## Controller Specifications

### Ports

- Dual 20Gbps or quad 10Gbps Ethernet

### Connectors

- 10GbE:
  - Two or four SFP+ ports (supporting 1G/10G) or;
  - Two or four RJ45 ports (with external 10GBASE-T PHY supporting 100M/1G/10G) or;
  - Two or four KR ports (supporting 1G/10G)
- 20GbE: two KR2 ports

### Certifications

- FCC A, UL, CE, VCCI, BSMI, C-Tick, KCC, TUV, and ICES-003

### Temperature

- Storage: less than 86°F (less than 30°C)

### Packaging

- 23mm × 23mm, 484-ball, flip-chip ball grid array with heat spreader (FCBGA-H)
- 1.0mm ball pitch

## Environmental/Equipment

### Compliance

- RoHS 6 compliant
- Halogen free

## Ordering Information

### Cavium FastLinQ BCM57840S, part number B57840SB1KFSBR

- Ships with a minimum order of 420 devices (60 devices per tray × 7 trays)



Follow us:     

[Corporate Headquarters](#) Cavium, Inc. 2315 N. First Street San Jose, CA 95131 408-943-7100

© 2015–2018 Cavium, Inc. All rights reserved worldwide. QLogic Corporation is a wholly owned subsidiary of Cavium, Inc. Cavium, QConvergeConsole, QLogic, and FastLink are registered trademarks or trademarks of Cavium, Inc. 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.