

QLogic[®] 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 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 for power savings with 10GBASE-T
- Virtualization environments
- 10GBASE-KR mezzanine card for server blades
- 10GBASE-T copper NIC with external 10G copper PHY (BCM84833, for example)

OVERVIEW

The QLogic[®] BCM57840S 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 BCM57840S includes quad-channel 10GBASE-KR and SFF-8431 for SFP+ 10Gb interfaces. The BCM57840S 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 utilization of 10GbE bandwidth, the QLogic BCM57840S offers QLogic switch-independent NIC partitioning (NPAR), which enables the segmentation of a single 10GbE 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 utilization while lowering infrastructure and operational costs.

The BCM57840S enables convergence of all possible network communications in a server, such as data network (LAN), FCoE storage network or block (for example, iSCSI). The BCM57840S 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 utilization, which results in significant system-level power savings.

The BCM57840S 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
- 2x10G, 4x10G, 2 × 20G modes

Host Interfaces

- PCI Express x8 3.0, 8GT/s compliant
- PCI Express x8 2.0, 5GT/s compliant
- PCI Express x8 1.1, 2.5GT/s compliant
- PCI Express lanes x1, x4, and x8
- No external dynamic random-access memory (DRAM) required; flow-through architecture
- PCI Express $\overline{\text{CLKREQ}}$ support
- SR-IOV
- Comprehensive IPv4 and IPv6 stateless offloads
- Broad OS and hypervisor support
- RSS and TSS
- Support for jumbo frames up to 9,600 bytes
- Network teaming, failover and load balancing
- MSI and MSI-X support
- Switch-independent NIC partitioning
- Generic routing encapsulation (NVGRE) packet task offloads
- Virtual extensible LAN (VXLAN) packet task offloads
- Generic Network Virtualization Encapsulation (Geneve) 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 functionality iSCSI initiator
- iSCSI boot and iSCSI crash dump support

FCoE

- Receiver and transmitter CRC offload
- Offloaded full Host Bus Adapter functionality 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
- Wake-on-LAN (WoL)

- 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 utilization 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 utilization for file-level storage protocols such as Common Internet File System (CIFS), Server Message Block Protocol (SMB), and NFS
 - Offloaded and accelerated iSCSI block storage with high I/O per second and low CPU utilization
- Accelerated FCoE
 - Offloaded and accelerated FCoE for Fibre Channel block storage with high I/O per second and low CPU utilization
- Performance-focused—Optimized for high throughput, low latency, and CPU utilization
 - Adaptive interrupt coalescing
 - Receive side scaling (RSS) reduces CPU utilization 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 functionality 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

- PCI Express 3.0 x8 (x8 physical connector)

Host Interrupts

- MSI-X supports independent queues

I/O Virtualization

- SR-IOV (128 maximum virtual functions per device)
- QLogic switch-independent NIC partitioning (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

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

- 1,500 bytes and larger (jumbo frames)

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 QLogic switch-independent partitioning modes
- VMware® NetQueue and Microsoft virtual machine queue (VMQ)

Compliance

- *IEEE 802.3ae (10Gb Ethernet)*
- *IEEE 802.1q (VLAN)*
- *IEEE 802.3ad (Link Aggregation)*
- *IEEE 802.3x (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 (Hardware Precision Time Protocol)*

Tools and Utilities

Management Tools and Device Utilities

- QConvergeConsole® (QCC) integrated network management utility graphical user interface (GUI) for Linux and Windows
- QLogic 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
- QLogic 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 <http://driverdownloads.qlogic.com>

Controller Specifications

Ports

- Dual 20Gbps or quad 10Gbps Ethernet

Connectors

- 10GbE: two or four SFP+ ports or RJ45 ports (with external 10GBASE-T PHY) or KR ports
- 20GbE: two KR ports

Certifications

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

Temperature

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

RoHS Compliance

- Green (RoHS 6 compliant and halogen free)

Packaging

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

Ordering Information

QLogic BCM57840S, part number B57840SB1KFSBR

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



Follow us:        Share:   

[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

© 2015-2017 QLogic Corporation. QLogic Corporation is a wholly owned subsidiary of Cavium, Inc. All rights reserved worldwide. FastLinQ, QConvergeConsole, QLogic and the QLogic logo are registered trademarks of QLogic Corporation. Microsoft is a registered trademark of Microsoft Corporation. PCI-SIG, PCI Express, and PCIe are registered trademarks of PCI-SIG Corporation. VMware and vSphere are trademarks or registered trademarks of VMware, Inc. All other brand and product names are trademarks or registered trademarks of their respective owners.

This document is provided for informational purposes only and may contain errors. QLogic reserves the right, without notice, to make changes to this document or in product design or specifications. QLogic 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 QLogic's future direction and intent are subject to change or withdrawal without notice and represent goals and objectives only.