




BCM57711 PRODUCT Brief



10-Gbps DUAL-PORT TCP, iSCSI, RDMA CONTROLLER WITH x8 LANE PCI EXPRESS® Ver. 2.0

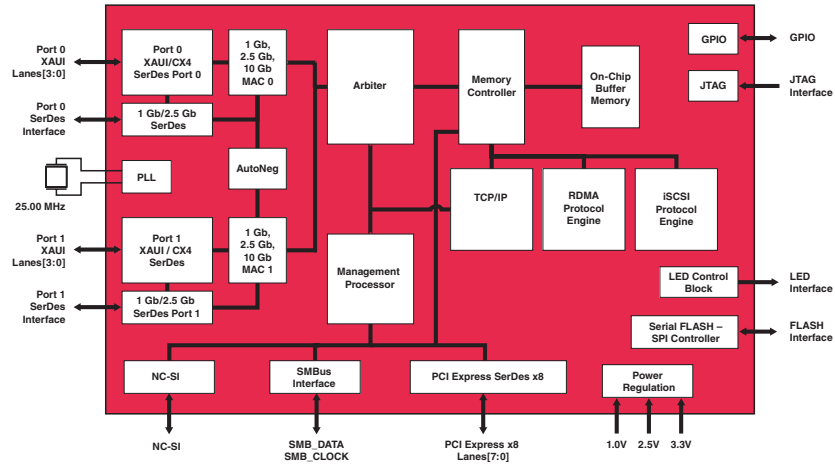
FEATURES

- **Single-chip solution for LAN on Motherboard (LOM) and converged network interface card (C-NIC) applications**
 - Integrated dual 10 Gbps MAC and dual XAUI™/10GBASE-CX4/10GBASE-KX4
 - Single 25.00 MHz clock crystal for 10 Gbps operation
 - Host interfaces
 - PCI Express x8 v2.0-compliant
 - Network interfaces
 - Dual-port 1000BASE-KX/2500BASE-KX/10GBASE-KX4/XAUI/10GBASE-CX4 interfaces for 1 Gbps/2.5 Gbps/10 Gbps operation
 - IEEE 802.3ap clause 73-compliant backplane operation
 - IEEE 802.3xx clause 37-compliant auto-negotiation for 1 Gbps and 2.5 Gbps
- **TCP**
 - Microsoft® TCP chimney compliant
 - Full FASTPATH® TCP for IPv4 and IPv6
 - No external memory required; flow-through architecture
- **iSCSI controller**
 - Off-loaded full HBA functionality iSCSI initiator
 - iSCSI boot and iSCSI crash dump support
- **RDMA controller (RNIC)**
 - Hardware-based data placement in application buffers without CPU intervention (for user and kernel modes)
 - Very low latencies
- **Simultaneous operation of Ethernet, TCP/IP, iSCSI, and RDMA modes**
- **Other performance features**
 - Receive side scaling (RSS) for IPv4 and IPv6
 - Teaming for L2, L4, and L5
 - Giant send offload (GSO) support
 - Jumbo frame support (9600 bytes)
 - TCP, IPv4, IPv6 checksum offload
 - TCP segmentation offload
 - Message signal interrupt (MSI/MSI-X) support
- **Robust manageability**
 - Network controller sideband interface (NC-SI)
 - PXE v2.1 remote boot
 - Wake-On-LAN (WOL)
 - SMASH-compliant
 - IPMI pass-through feature
 - Statistics gathering (SNMP MIB II, Ethernet MIB (802.2x, clause 30))
 - Comprehensive diagnostic and configuration software suite
 - ACPI v2.0b-compliant power management
- **Advanced network features**
 - Virtual LANs—802.1q VLAN tagging
 - 802.3x flow control
 - Congestion management support

SUMMARY OF BENEFITS

- **Industry's first 1000 Mbit/2500 Mbit/10 Gbps SerDes-based TCP/IP solution—power and space optimized for server blade, rack tower, and C-NIC applications**
 - Four-lane XAUI for 10 Gbps operation on server blade backplane or rack/tower using optical or CX-4 cables
 - Single-lane 1 Gbps or 2.5 Gbps for server blade with one lane
 - 1-Gbps operation over copper or optical fiber for rack and tower servers
- **Extremely low CPU utilization for TCP/IP applications**
 - 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 CIFS/SMB and NFS
 - Offloaded and accelerated iSCSI block storage with high I/O per second and low CPU utilization
- **High-performance clustered systems with low latency for latency-sensitive applications (e.g., MPI-based applications)**
- **Future-proof**
 - Firmware-based flexible implementation for TCP, RDMA, and iSCSI can accommodate specification changes and interoperability issues
- **Shares existing software base with second generation NetXtreme® II controller family of products**
- **Performance-focused—optimized for high throughput lowest latency and CPU utilization**
 - Adaptive interrupt coalescing
 - 2.5 Gbps Ethernet
 - 10 Gbps Ethernet over XAUI/CX4/IEEE802.3ap
 - RSS reduces CPU utilization on multi-CPU systems
 - MSI/MSI-X allows interrupt distribution in a multi-CPU system
 - PCI Express host interface allows for low-latency access to CPU and host memory resources
- **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, WOL
 - IPMI pass-through allows on-board management controllers access to the network in OS-present and OS-absent states
- **Server class reliability, availability, and performance features**
 - Link aggregation and load balancing
 - Switch-dependent
 - 802.3ad (LACP), generic trunking (GEC/FEC)
 - Switch- and NIC-independent
- **Direct support of CX4 Interface for LOM and C-NIC applications**
- **Glue-less interconnect**
 - Glue-less interconnect for 10GBASE-KX4 10 Gbps backplane interface for blade servers
 - Glue-less interconnect to BCM8704/5 device for optical XFP solution
- **Direct connection to XENPAK or X2 10-Gbps optics modules**
- **Configurable LED interface**
- **676-pin 27 mm x 27 mm PBGA+HS package; RoHS-compliant**
- **1.0V core voltage; 3.3V tolerant I/Os**
- **IEEE 1149-1 compliant JTAG support**

OVERVIEW



BCM57711 Block Diagram

The BCM57711 provides a fully integrated Layer 4 and Layer 5 solution—TCP/IP, RDMA, and iSCSI 1.0 along with dual 1000/2500/10GBASE-X Ethernet, IEEE 802.3-compliant media access control (MAC), and physical layer transceiver for high-performance network applications.

By itself, the BCM57711 provides a complete single-chip dual-channel 10GBASE-KX4/10GBASE-CX4 C-NIC with a TCP/IP, RDMA NIC (RNIC), and iSCSI 1.0 HBA functionality for NIC or LOM solutions.

Interfacing to various 10 Gbps SerDes products from Broadcom, the BCM57711 can be used in multiple application scenarios. Using Broadcom BCM87XX devices, the XFI interface allows interfacing to XFP modules. Various companion devices can be used for 10GBASE-KR or other copper interfaces.

The BCM57711 is a third-generation C-NIC and different from Layer-2 Ethernet network controllers. When transmitting, the BCM57711 can access data posted in application buffers or other buffers and process the TCP/IP, iSCSI, and RDMA layers, add the appropriate headers and fields (e.g., CRC and markers), and transmit the data, thereby relieving the host CPU from these time-consuming operations. On the receive path, the BCM57711 processes the frame up to the highest layer present for TCP, iSCSI, or RDMA strip headers, and special fields (e.g., CRC and markers) and places the data with zero copy in the best buffer available for the data—an application buffer, TCP level buffer, or a generic buffer.

With the appropriate configuration, the BCM57711 can simultaneously support the following functions:

- RDMA network interface controller (RNIC)
- iSCSI host bus adapter (HBA)

- TCP/IP-enabled network accelerator
- Ethernet L2 network interface controller
- Management

Target Applications of the BCM57711

- Converged NIC (C-NIC), LAN on Motherboard (LOM), blade tower, and rack server
- 10 Gbps dual XAUI/CX4 C-NIC/LOM
- 10 Gbps dual XAUI/KX4 blade server
 - Autodetect and autonegotiate to 1000/2500BASE-X
- 2.5 Gbps Ethernet server blade
- iSCSI v1.0 host bus adapter (HBA)
- RDMA-enabled network interface card (RNIC)

<i>Converged Network Interface Card (C-NIC) Designs</i>	
Dual 1000BASE-X	PCI Express x4 or x8
Dual 2500BASE-X	PCI Express x4 or x8
Dual 10 Gbps CX4	PCI Express x8
Dual 10 Gbps XFP using BCM87xx	PCI Express x8
10G XFP, 10GBASE-T (via companion device)	PCI Express x8

<i>LAN on Motherboard (LOM) Designs</i>	
Dual 1000BASE-X	PCI Express x4 or x8
Dual 2500BASE-X	PCI Express x4 or x8
10G Mezzanine Card for Server Blades	PCI Express x8
10GBASE-CX4, 10GBASE-KX4, 10GBASE-KR (via companion device), 10GBASE-T (via companion device)	PCI Express x8

Broadcom®, the pulse logo, Connecting everything®, the Connecting everything logo, FASTPATH®, and NetXtreme® II are among the trademarks of Broadcom Corporation and/or its affiliates in the United States, certain other countries and/or the EU. Any other trademarks or trade names mentioned are the property of their respective owners.

Connecting
everything®



BROADCOM CORPORATION

5300 California Avenue
Irvine, California 92617

© 2009 by BROADCOM CORPORATION. All rights reserved.

57711-PB00-R 10/08/09

Phone: 949-926-5000
Fax: 949-926-5203
E-mail: info@broadcom.com
Web: www.broadcom.com