

## Product Brief



### Highlights

- Single, dual, or quad-port with 400Gb/s aggregate throughput
- 400/200/100/50/25GbE
- 8x SerDes: 100/50G PAM-4 or 10/25G NRZ
- PCIe 5.0 x16, multihost up to four hosts
- Up to 5 meter DAC cable support
- RoCEv2 with DCN and SARA
- TLS/QUIC inline encryption
- Advanced timing support: PTP, SyncE
- Transmit pacing
- HW secure boot and attestation
- TruFlow™ 400G flow processing offload
- Peer memory direct
- PCIe embedded analyzer (PEA)
- Comprehensive stateless offloads
- Standards-based system management

# BCM57608

## Quad-Port 400Gb/s Ethernet Controller with PCIe 5.0 and 100G PAM-4 SerDes

### Description

Building upon the success of the widely deployed Broadcom® NIC architecture, the BCM57608 series Ethernet controllers continue to drive innovation by supporting 400Gb/s aggregate throughput while maintaining the industry's lowest power and best thermal solution. The BCM57608 supports network interfaces of 400/200/100/50/25GbE.

Incorporating industry-leading PHY technology, the BCM57608 provides highly efficient network connectivity using eight lanes of built-in SerDes, each supporting 100/50G PAM-4 and 10/25G NRZ. When paired with the Broadcom market-leading StrataXGS® or StrataDNX™ Ethernet switches, the BCM57608 ensures the highest level of performance, interoperability, and features. This includes the longest reach for passive direct attached cables (DAC) of up to 5 meters, to accommodate most in-rack and rack-to-rack deployments.

The BCM57608 supports the fourth generation, standards-based RDMA over converged Ethernet (RoCE) with hardware-based congestion control. Broadcom RoCE congestion control not only delivers the lowest latency in real-life scenarios, but also dramatically reduces the complexity of RoCE deployment at scale.

The BCM57608 fourth-generation 400G TruFlow™ engine integrates hardware acceleration with enhanced programmability for the parser/lookup/match/action engine to allow for rapid implementation of new flow types. TruFlow increases virtual machine density and improves application performance.

The BCM57608 supports technology-leading security, enabling the industry's most secure server platform with Broadcom HW secure boot (ROT) and attestation support.

## Applications

- Cloud and enterprise data center servers
- Artificial intelligence (AI) and machine learning (ML) clusters
- NVMe storage disaggregation
- 5G wireless RAN
- Network function virtualization
- Mobile edge computing
- High performance computing (HPC)

## Benefits

The BCM57608 leverages industry-leading technology to deliver best-in-class results:

- Lowers data center total cost of ownership (TCO) with the industry's lowest-power, 400G-capable NIC
- Unleashes the networking performance of PCIe 5.0-capable servers
- Best-in-class 100G SerDes increases cable reach and reduces power cost over optical
- Enables large-scale RoCE deployment with switch aware rate adjustment (SARA) and drop congestion notification (DCN)
- Inline encryption support to offload CPU and improve system and application performance
- Advanced timing support to enable precision timing with sub-nanosecond timing
- Improves network reliability and real-time control with in-band telemetry
- Accelerates applications with kernel-bypass: RDMA, DPDK, and SR-IOV
- Improves CPU efficiency for I/O with comprehensive stateless offloads
- Enhances system and network management

## Key Features

### Network Interface:

- 8 SerDes capable of 100/50G PAM4 and 10/25G NRZ
- Support for up to four ports
- 1x 25/50/100/200/400GbE
- 2x 25/50/100/200GbE
- 4x 25/50/100GbE
- 400Gb/s total bandwidth
- Auto-negotiation with auto-detect

### Host Interface:

- 16 lanes of PCI Express 5.0
- Link rates: 32, 16, 8, 5, 2.5 GT/s
- Lane configuration: x16, x8, x4, x2, and x1
- PCIe bifurcation
- Multi-host / multi-root up to four
- MSI-X support

### Security:

- Data Security:
  - kTLS hardware offload encryption/decryption support
  - QUIC hardware offload encryption/decryption support
- Platform Security:
  - Hardware secure boot (RoT)
  - Attestation (SPDM 1.2)
  - OCP secure recovery
  - Secure wipe and restore
  - OCP Silver security badge (certification pending)

## Key Features (cont.)

### Networking and Virtualization:

- RoCEv2
- Multi-Queue, NetQueue, and VMQ
- Single root I/O virtualization
- VF isolation and protection
- VXLAN, GRE, NVGRE, Geneve, and IP-in-IP
- Tunnel-aware stateless offloads
- Edge virtual bridging (EVB)
- Stateless TCP offloads: IP/TCP/UDP checksum, LSO, LRO, GRO, TSS, RSS, aRFS, interrupt coalescing

### TruFlow™ Flow Processing:

- Flexible matching key
- NAT and NATP
- Tunnel encap/decap
- Custom tunnel processing
- Connection tracking
- Flow aging
- Sampling and mirroring
- Rate-limiting and metering
- Flow-based statistics
- Network traffic hairpin

### RoCEv2

- Standards-based
- DCQCN
- Peer memory direct
- DCN and SARA
- Automated configuration

### Timing and Synchronization

- IEEE 1588v2
- Class C/D timing
- 1-ns time-stamping accuracy
- SyncE
- PPS
- Linux SO\_TXTIME
- Pacing: fair queuing offload (FQO)
- O-RAN LLS-C1/2/3

### Manageability

- Network controller sideband interface (NC-SI)
- Management Component Transport Protocol (MCTP)
- MCTP over SMBus/I2C
- MCTP over PCIe VDM
- NC-SI over MCTP
- Platform Level Data Model (PLDM): base, monitoring/control, and firmware update
- PLDM over MCTP
- I2C support for device control and configuration

### Network Boot

- PXE boot
- iSCSI boot
- UEFI support for x86 and Arm servers

## Ordering Information

Part Number	Throughput	Host Interface	Network Interface	Package
BCM57608B1KF5BG	400 Gb/s	×16 PCIe 5.0	1 × 400GbE, 2 × 200GbE, or 4 × 100GbE	23 mm x 23 mm FCGBA