Marvell® FastLinQ® 45000 Series

Industry Leading 10/25/40/50/100GbE RoCE, iWARP, FCoE, iSCSI, and SR-IOV PCIe® 3.0 Ethernet Controllers

- Quad-port 10GbE and 25GbE applications
- Dual-port 25GbE, 40GbE, and 50GbE applications
- Single-port 100GbE applications
- Secure firmware update process with private/public key encryption technology prevents hackers from altering adapter
- VXLAN, NVGRE, GRE, and GENEVE tunneling offloads
- Universal RDMA—RoCE, RoCEv2, and iWARP
- NVMe-oF NVMe® over Fabrics with Universal RDMA and TCP
- T10 Protection Information

Marvell FastLinQ 45000 Series Ethernet Controllers are a fifth-generation solution designed for high-volume, converged network applications. The QL45604 supports speeds of 100Gbps, 50Gbps, 40Gbps, 25Gbps, 10Gbps, and 1Gbps. The QL45204 supports speeds of 25Gbps, 10Gbps, and 1Gbps. Marvell FastLinQ 45000 Series Controllers enable single root I/O virtualization (SR-IOV), RDMA over converged Ethernet (RoCE), Internet wide area RMDA protocol (iWARP), iSCSI, Fibre Channel over Ethernet (FCoE), and data center bridging (DCB). They also support 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 Marvell FastLinQ 45000 Series is a complete solution that enables leading-edge features for the enterprise and cloud (independent of server form-factor), while significantly raising the performance bar. This solution enables stateful and stateless offloads and includes advanced features such as network virtualization offload, storage offloads, secure firmware update with private/public key encryption, and FastLinQ SmartAN™ for simplified connectivity to switches without user intervention.

The QL45604 includes support for single-port 100GBASE, dual-port 50GBASE, dual-port 40GBASE, quad-port 25GBASE, and quad-port 10GBASE applications. The QL45204 includes support for up to quad-port 25GBASE Marvell FastLinQ 45000 Series Controllers integrate four IEEE 802.3-compliant MACs and support the network controller-sideband interface (NC-SI). Host-to-baseboard management controller (BMC) communication is also supported on top of the NC-SI to permit high-speed communication between the local host and the BMC or management controller (MC).

Marvell FastLinQ 45000 Series Controllers enable networked storage using block-based storage (iSCSI or FCoE) and file-based storage (CIFS or NFS). Clustering and interprocess communications (IPC) are supported with RoCE and iWARP. They can simultaneously support all offload traffic types on each of the ports. Offloading results in superior storage and networking performance, as well as low CPU utilization, which in turn results in significant system-level power savings.

Marvell FastLinQ 45000 Series Controllers are designed for PCIe 3.0 and are also compatible with the *PCI Express Base Specification*, revisions 2.0 and 1.1. PCIe supports MSI and MSI-X capabilities. Each port supports multiple physical functions.

Marvell FastLinQ 45000 Series Controllers support IEEE 1588 precision timing protocol (PTP) and IEEE 802.1AS, providing a method of synchronization between *master* and *slave* clocks over a LAN.

Features

The FastLinQ 45000 Series Controllers include the following features:

Network Interfaces

- Blade and dense servers:
  - 1000BASE-KX
  - 10GBASE-KR
  - 25GBASE-KR-S
  - 40GBASE-KR4
  - 50GBASE-KR2
  - 100GBASE-KR4
- Rack, tower, and dense servers:
  - SFF8431 Annex E 10GbE (direct attach copper)
  - 1000BASE-CX/SX/LX
  - 10GBASE-T (with external 10GBASE-T PHY)
  - 10GBASE-SR/LR
  - 25GBASE-CR-S
  - 40GBASE-SR4 and 40GBASE-CR4 (direct attach copper)
  - 50BASE-CR2
  - 100GBASE-CR4/SR4/LR4

iSCSI Controller

- Offloaded full Host Bus Adapter (HBA) functionality iSCSI initiator
- iSCSI boot and iSCSI crash dump support

FCoE

- Offloaded full HBA functionality FCoE initiator
- FCoE boot

Universal RDMA (RoCE/iWARP)

- Hardware-based data placement in application buffers without CPU intervention (for user and kernel modes)
- Low latency

Data Integrity

- ECC and byte parity protection
- T-10 CRC

Robust Manageability

- Simplifies deployment and trouble-shooting using QLogic® Control Suite (QCS) CLI, QConvergeConsole® (QCC) PowerKit, Unified Extensible Firmware Interface (UEFI) human interface infrastructure (HII), in-OS utilities, QCC Center GUI and ESXCLI Plug-ins

DCB

- Enhanced Transmission Selection (ETS) (IEEE 802.1Qaz)
- Quantized Congestion Notification (QCN)-capable (IEEE 802.1Qau)
- Priority-based Flow Control (PFC) (IEEE 802.1Qbb)
- Up to four traffic classes

Forward Error Correction (FEC)

- FireCode BASE-R IEEE 802.3-2018 Clause 74 (FC-FEC) on 10G through 100G interfaces
- Reed Solomon IEEE 802.3-2018 Clause 91 (RS-FEC) on 100G interfaces

Host Interfaces

- PCIe 3.0 x16 (8GTps), 2.0 (5GTps), and 1.1 (2.5GTps)

Benefits

Each FastLinQ 45000 Series Controller is an SR-IOV converged solution—power and space is optimized for blade, rack, and tower servers, and Converged Network Adapter (CNA) applications.

FastLinQ 45000 Series Controllers provide the following benefits:

Extremely Low CPU Utilization for iSCSI, FCoE, and RDMA 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 IOPS and low CPU utilization

Accelerated FCoE

- Offloaded and accelerated FCoE for Fibre Channel block storage with high IOPS 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 allow interrupt distribution in a multi-CPU system

Robust and Highly Manageable

- NC-SI, Serial gigabit media independent interface (SGMII) for 1Gb BMC interconnect, Management Component Transport Protocol (MCTP) over PCIe vendor-defined messages (VDM), and MCTP over system management bus (SMB) system management functionality over shared infrastructure
- Reliable delivery of management traffic
- Pre-execution environment (PXE) v2.1, advanced configuration and power interface (ACPI) v2.0b, Wake-on-LAN (WoL)
- Host-to-BMC communication for connectivity between the local host and the MC or BMC

Server-class Reliability, Availability, and Performance Features

- Network teaming, failover, and load balancing
  - Switch independent NIC teaming/bonding
  - Switch dependent NIC teaming/bonding such as link aggregation control protocol (IEEE 802.3ad LACP) and generic trunking
- Switch-independent NIC partitioning (NPAR) with up to 8 partition assignments per adapter, and NIC extended partitioning (NPAReP) with up to 16 partition assignments per adapter, both concurrently available with SR-IOV
- Marvell Data Plane Development Kit (DPDK) high-speed packet processing engine
- Marvell Flow Filtering is supported on Linux® using the `ethtool -u/-U` commands. See the See the n-tuple Flow Filtering and Steering FastLinQ 41000/45000 Series Adapters Deployment Guide for more information
- Comprehensive stateless offloads
- Energy Efficient Ethernet (EEE) support for reduced idle power consumption in RJ-45-based networks (10GBASE-T variants only)

PACKAGING

- 27mm × 27mm; 676 pins

Ordering Information

- QL45604-B0G
  - Supports all speeds: 1Gbps through 100Gbps
- QL45204-B0G
  - Supports 1Gbps through 25Gbps

Note: All advertised features are enabled in the hardware. Actual feature availability is dependent on software driver releases. See the release notes.

MARVELL

To deliver the data infrastructure technology that connects the world, we're building solutions on the most powerful foundation: our partnerships with our customers. Trusted by the world's leading technology companies for 25 years, we move, store, process and secure the world's data with semiconductor solutions designed for our customers' current needs and future ambitions. Through a process of deep collaboration and transparency, we're ultimately changing the way tomorrow's enterprise, cloud, automotive, and carrier architectures transform—for the better.