

Marvell® 41000 Series

Industry Leading 25/10GbE Controllers

Overview

FastLinQ® 41000 Series Ethernet Controllers are single, dual, and quad-port Ethernet solution designed for high-volume, converged network applications. The QL41000 Series Controllers support speeds of 25/10/1Gbps and enable single root I/O virtualization (SR-IOV), universal Remote Direct Memory Access (RDMA) over converged Ethernet (RoCE), iSCSI, Fibre Channel over Ethernet (FCoE), and data center bridging (DCB). They also support PCI Express® (PCIe®) Gen 3, along with embedded virtual bridging and other switching technologies for virtual machine (VM)-to-VM switching.

The FastLinQ 41000 Series is a complete solution that enables leading-edge features for the enterprise and cloud (independent of server form factor), while significantly accelerating network performance. This controller 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 QL41000 Series Controllers include support for 25GbE applications (single port and dual port) and 10GbE applications (single port, dual port, and quad port). CNAs support iSCSI and FCoE hardware-based offload. FastLinQ 41000 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 NC-SI to permit high-speed communication between the local host and the BMC or management controller (MC).



- Single, dual, and quad-port 10GbE applications
- Single and dual-port 25GbE applications
- Hardware-based iSCSI and FCoE offloads
- Secure firmware update process with private/public key encryption technology prevents hackers from altering adapter
- Low-power, single-chip solution with adaptive voltage scaling
- Universal RDMA—RoCE, RoCEv2, and iWARP
- VXLAN, NVGRE, GRE, and GENEVE tunneling offloads
- T10 Protection Information
- SR-IOV and NPAR
- NVMe-oF NVMe™ over Fabrics with RDMA and TCP

Features

- PCI Express® (PCIe®) Gen 3 x8 (8GT/s) support
- Full line-rate performance across all ports
- Broad operating system (OS) and hypervisor support
- Network boot support
 - iSCSI remote boot
 - FCoE remote boot from LUN
 - Preboot Execution Environment (PXE) 2.0
 - Unified Extensible Firmware Interface (UEFI) support
- Simplifies deployment and troubleshooting using QConvergeConsole® (QCC) GUI, QLogic® Control Suite (QCS) CLI, QCC PowerKit, UEFI human interface infrastructure (HII), in-OS utilities, QCC vCenter GUI and ESXCLI Plug-ins, and OpenStack® integration
- 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
- Marvell Data Plane Development Kit (DPDK) highspeed packet processing engine delivers up to 38 million packets per second at 64B frame sizes
- Marvell Flow Filtering is supported on Linux® using the ethtool -u/-U commands. See the [n-tuple Flow Filtering and Steering FastLinQ 41000/45000 Series Adapters Deployment Guide](#) for more information.
- Universal RDMA technologies—RDMA over Converged Ethernet (RoCE), RoCEv2, and Internet wide area RDMA protocol (iWARP)
- Energy Efficient Ethernet (EEE) support for reduced idle power consumption in RJ-45-based networks (10GBASE-T variants only)
- MSI and MSI-X support
- IPv4 and IPv6 stateless offloads
- PCI-SIG® single root input/output virtualization (SR-IOV) with up to 192 virtual functions
- Comprehensive stateless offloads
- Auto negotiation: 1G/10G (10GBASE-T only)
- FastLinQ SmartAN for simplified connectivity with 1/10/25G interfaced switches. (SFP interfaces can accept direct attach cable (DAC) or optical (discrete or active optic cable (AOC)) connections).
- RX/TX multiqueue
 - VMware® NetQueue
- RX/TX multiqueue (continued)
 - Windows® Hyper-V® Virtual Machine Queue
 - Linux Multiqueue
- Tunneling offloads
 - Windows Network Virtualization using Generic Routing Encapsulation (NVGRE)
 - Linux Generic Routing Encapsulation (GRE)
 - VMware, Windows, and Linux Virtual Extensible LAN (VXLAN)
 - Linux and VMware Generic Network Virtualization Encapsulation (GENEVE)
- Receive side scaling (RSS)
- Transmit side scaling (TSS)
- Support for virtual LAN (vLAN) tagging
- Support for jumbo frames larger than 1,500 bytes (up to 9,600 bytes)
- Network teaming, failover, and load balancing
 - Switch independent NIC teaming/bonding
 - Switch dependent NIC teaming/bonding such as link aggregation control protocol (LACP) and generic trunking
- Data center bridging (DCB)
 - Priority-based Flow Control (PFC) over
- VLAN's 3-bit Priority Code Point (PCP) field
 - PFC over the IP header's 3-bit Differentiated Services Code Point (DSCP) field
 - Enhanced Transmission Selection (ETS)
 - Explicit Congestion Notification (ECN or CN)
 - Data Center Quantized Congestion Notification (DCQCN)
- Non-offloaded Storage over Ethernet
 - iSCSI using OS-based software initiators
- Offloaded storage over Ethernet
 - Increases server performance with full hardware offload for storage traffic
 - Industry-leading FCoE-Offload performance of up to 3.6 million IOPS, suitable for high-density server virtualization and large databases
 - Industry-leading iSCSI-Offload performance of up to 2.9 million IOPS, suitable for a diverse set of applications leveraging the flexibility of iSCSI

Key Features

Network Interfaces

Blade and Dense Servers

- 1000BASE-KX
- 10GBASE-KR
- 25GBASE-KR

Rack, Tower, and Dense Servers

- 1000BASE-CX
- 1000BASE-SX
- 1000BASE-LX
- 1000BASE-T (with external 1GBASE-T PHY)
- SF8431 Annex E 10GbE (direct attach copper)
- 10GBASE-SR
- 10GBASE-LR
- 10GBASE-T (with external 10GBASE-T PHY)
- 25GBASE-CR
- 25GBASE-SR
- 25GBASE-LR

Host Bus Interface

Bus Interface

- PCI Express (PCIe) Gen 3 x8 (x8 physical connector)

Host Interrupts

- MSI-X supports independent queues

I/O Virtualization and Multitenancy

- SR-IOV (up to 192 virtual functions)
- Switch-independent NPAR, NPAReP
- GRE and NVGRE packet task offloads
- VXLAN packet task offloads
- GENEVE packet task offloads

Compliance

- PCI Base Specification, rev. 3.1
- PCI Bus Power Management Interface Specification, rev. 1.2
- Advanced configuration and power interface (ACPI) v2.0

I/O Virtualization

- SR-IOV
- NetQueue
- Dynamic virtual machine queue (DVMQ)
- Multiqueue support
- Switch-dependent and -independent NIC partitioning (NPAR)

Tunneling Offloads

- VXLAN
- NVGRE
- GENEVE
- GRE

Network, Teaming, Failover, and Load Balancing

- Switch independent
- Switch dependent (IEEE 802.3as Link Aggregation Control Protocol (LACP) and generic trunking (GEC/FEC))

Throughput

- 10/25Gbps line rate for single and dual port
- 1/10GBASE-T Auto Negotiation

Ethernet Frame

- 1,500 bytes and larger (jumbo frame)

Stateless Offload

- TCP segmentation offload (TSO)
- Large send offload (LSO)
- VMware large receive offload (LRO)
- Linux generic receive offload (GRO)
- Generic segmentation offload (GSO)
- TCP and user datagram protocol (UDP) checksum offloads
- Receive segment coalescing (RSC)
- Interrupt coalescing
- RSS and TSS
- VMware NetQueue, Microsoft Hyper-V VMQ (up to 208 dynamic queues)/Virtual Machine Multi-Queue (VMMQ)/Virtual Switch RSS (vRSS), Linux Multiqueue, and Virtual Machine Device queues (VMDq)
- DPDK
- Universal RDMA

Compliance

- IEEE Specifications
 - 802.1AS (Precise Synchronization)
 - 802.1ax-2008 (Link Aggregation)
 - 802.1p (Priority Encoding)
 - 802.1q (VLAN)
 - 802.1Qau (CN)
 - 802.1Qaz (DCBX and ETS)
 - 802.1Qbb (PFC)
 - 802.3-2018 Annex 31B (Ethernet Pause Flow Control)
 - (RJ45) 802.3-2018 Clause 78 EEE (Energy Efficient Ethernet)
 - 802.3by-2016 (25G Ethernet)
 - (25GbE) 802.3-2018 Clause 110 (Direct Attach Copper), Clause 112 (SR optical), and Clause 114 (LR optical) (25G Ethernet)
 - (10GbE SFP+) 802.3-2018 Clause 52 (10Gb Ethernet Optical)
 - 1588-2002 PTPv1 (Precision Time Protocol)
 - 1588-2008 PTPv2
 - (10GbE SFP+) SFF8431 Annex E (10Gb Direct Attach Copper)
- RFCs
 - IPv4 (RFC 791)
 - IPv6 (RFC 2460)

Board Firmware Features

- Secure Firmware Update process
- Smart Auto Negotiation (FastLinQ SmartAN)

Compliance

- RoHS compliant

Data Integrity

- ECC and byte parity protection
- T-10 CRC

RDMA

Universal RDMA

- RoCE
- RoCEv2
- iWARP
- Storage over RDMA: iSER, SMB Direct, S2D, and NVMe-oF
- NFSoRDMA
- iSCSI extensions for RDMA (iSER)
- Low latency

RDMA Use Cases

- S2D
- PVRDMA
- Live Migration
- SMB Direct
- NVMe-oF
- NFS
- CEPHS over RDMA

FCoE

- Offloaded full HBA function FCoE initiator
- Performance of up to 3.6 million FCoE IOPS
- Technical Committee T11 in the Fibre Channel
 - Link Services (FC-LS) Specification's N_Port ID virtualization (NPiV) on Linux, Windows, and VMware

iSCSI

- Offloaded full HBA function iSCSI initiator
- Performance of up to 2.9 million iSCSI IOPS
- iSCSI crash dump support

Forward Error Correction (FEC)

- FireCode or BASE-R IEEE 802.3-2015 Clause 74 or FC-FEC
- Reed Solomon IEEE 802.3by Clause 91 or RS-FEC

Tools and Utilities

Management Tools and Device Utilities

- QCS Command Line Interface (CLI) for Linux and Windows
- QCC integrated network management utility (GUI) for Linux and Windows
- QCC Plug-in for vSphere (GUI), and ESXCLI plug-in for VMware
- QCC PowerKit (Windows PowerShell®) cmdlets for Linux, VMware, and Windows
- Pre-boot UEFI HII system BIOS device configuration pages
- Native OS management tools for networking
- Wake-on-LAN (WoL)
- Management Interfaces
 - NC-SI
 - Management Component Transport Protocol (MCTP) over PCIe vendor-defined messages (VDM)
 - MCTP over system management bus (SMB)
 - SMD and reduced media-independent interface (RMII)
 - Serial gigabit media independent interface (SGMII) for 1Gb BMC interconnect
 - Host-to-BMC communication

Boot Support

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

Operating System Support

- For the latest applicable operating system information, see www.marvell.com

DCB

- Enhanced Transmission Selection (ETS) (IEEE 802.1Qaz)
- Priority-based Flow Control (PFC) (IEEE 802.1Qbb)
- Up to four traffic classes
- Data Center Quantized Congestion Notification (DCQCN)/Explicit Congestion Notification (ECN) (IEEE 802.1Qau)

Host Bus Interface Specifications

Bus Interface

- PCIe 3.1 x8 (8GTps), 2.1 (5GTps), and 1.1 (2.5GTps)

Packaging

- 19mm × 19mm
- 525 pins
- Ball pitch: 0.8mm

41000 Series AVS¹ Controllers Ordering Information

| Part Number | Description |
|--------------|---------------------------|
| QL41101A-A2G | Single-port 10GbE |
| QL41102A-A2G | Dual-port 10GbE |
| QL41104A-A2G | Quad-port 10GbE |
| QL41201A-A2G | Single-port 25GbE/10GbE |
| QL41202A-A2G | Dual-port 25GbE/10GbE |
| QL41262A-A2G | Dual-port 25GbE/10GbE CNA |
| QL41162A-A2G | Dual-port 10GbE |

1 AVS = adaptive voltage scaling



Marvell first revolutionized the digital storage industry by moving information at speeds never thought possible. Today, that same breakthrough innovation remains at the heart of the company's storage, networking and connectivity solutions. With leading intellectual property and deep system-level knowledge, Marvell semiconductor solutions continue to transform the enterprise, cloud, automotive, industrial, and consumer markets. For more information, visit www.marvell.com.

© 2020 Marvell. All rights reserved. The MARVELL mark and M logo are registered and/or common law trademarks of Marvell and/or its Affiliates in the US and/or other countries. This document may also contain other registered or common law trademarks of Marvell and/or its Affiliates.

Marvell_41000 Series_PB Revised: 05/20