Marvell® FastLinQ® QL41162HLRJ-CI

Multiport 10Gb Converged Network Adapter with Universal RDMA

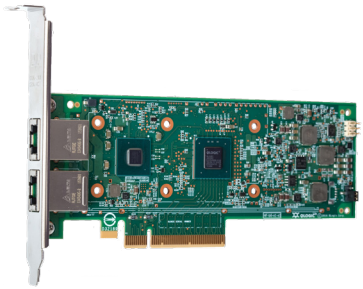

- Delivers full line-rate 10GbE performance
- Universal RDMA—Delivers the choice and flexibility with concurrent support for RoCE, RoCEv2, and iWARP technologies
- Secure firmware update process with private/public key encryption technology prevents hackers from altering adapter
- Enables provisioning of multiple QoS backed Ethernet functions for greater deployment flexibility through switch-independent NIC partitioning
- Boosts host CPU efficiency with hardware offloads for GRE, NVGRE, GENEVE, and VXLAN tunnels
- Low-cost and easy-to-install RJ45 connectivity that is compatible with existing 1GbE

The FastLinQ QL41162HLRJ-CI Converged Network Adapter (CNA) with Universal Remote Direct Memory Access (RDMA)—available in 10GBASE-T (RJ45)—supports LAN (TCP/IP) traffic at 10GbE line-rate speeds. The QL41162HLRJ-CI provides extremely low host CPU usage by enabling full stateless offloads to meet the performance requirements of the most demanding enterprise applications.

The QL41162HLRJ-CI leverages Marvell's 15+ years of expertise in Ethernet by providing the highest levels of performance, efficiency, and scalability for the enterprise data center.

For more effective use of the 10GbE bandwidth, the QL41162HLRJ-CI Converged Network Adapter offers switch-independent NIC partitioning (NPAR), which enables segmentation of a single 10GbE port into multiple network partitions and dynamic allocation of bandwidth to each port. The segmentation allows IT organizations to optimize resources while lowering infrastructure and operational costs.

The evolution of data centers—triggered by high-density server virtualization, software-defined networking (SDN), and multitenant cloud computing platforms—demands a high-performance 10GbE solution that boosts CPU efficiency and reduces capital expenditures (CAPEX) and operational expenditures (OPEX) of the migration to 10GbE. The QL41162HLRJ-CI is the best choice for workload-intensive computing environments, providing reliable, high-performance 10GbE connectivity solutions.

FEATURES

- 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
- 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)
- IPv4 and IPv6 stateless offloads
- Comprehensive stateless offloads

Key Features

- General Specifications
  - Ports: 2
  - Port Speeds: 1, 10
  - Connectors: BASE-T (supports Auto-Negotiation, 1GbE full duplex, and 10GbE full duplex)
  - Form Factor: PCIe MD2
  - Media: RJ-45CAT
  - Supports IEEE 802.3az (EEE)
- Storage
  - Universal RDMA (RoCE/ RoCEv2/iWARP)
  - NVMe-oF over TCP/RDMA
  - FCoE-Offload
  - iSCSI-Offload
- Virtualization and Cloud
  - Concurrent SR-IOV/NPAR
  - DPDK
  - Flow Filtering
  - Tunning Offload (VXLAN/ GENEVE/NVGRE/GRE)
- Physical Specs
  - Cooling Requirements: 150/55 LFM/ °C

- 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
  - 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 NPAR with up to 16 partition assignments per adapter
- Marvell Data Plane Development Kit (DPDK) high-speed 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
- MSI and MSI-X support
- PCI-SIG® single root input/output virtualization (SR-IOV) with up to 192 virtual functions
- Auto negotiation: 1G/10G (BASE-T)
- RX/TX multiqueue
  - VMware® NetQueue
  - Windows® Hyper-V® Virtual Machine Queue
  - Linux Multiqueue
- 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)
  - Data center bridging capability exchange protocol (DCBX) link layer discovery protocol (LLDP)
  - Priority-based flow control (PFC)
  - Traffic Class over VLAN's 3-bit priority code point (PCP) field or Traffic Class 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

BENEFITS

Simplified Migration to 10GbE

The Marvell FastLinQ QL41162HLRJ-CI Adapter features a high-speed, flexible architecture and switch-independent NPAR technology. Designed for both physical and virtual environments, this switch-agnostic approach enables administrators to split up the 10GbE network pipe to divide and reallocate bandwidth and resources, as needed, at the adapter level.

- Customers deploying rack and tower servers with multiple GbE adapters can greatly benefit from consolidating multiple network adapters and freeing up PCI slots for other add-in card upgrades.
- With NPAR, the QL41162HLRJ-CI Adapter can further partition its network bandwidth into multiple virtual connections, making 1 adapter appear as 16 adapters to the OS for use by the applications.
- NPAR greatly simplifies the physical connectivity to the server, reduces implementation time, and lowers the acquisition cost of 10GbE migration.
- Available in 10GBASE-T, the QL41162HLRJ-CI Adapter is the ideal choice for migrating multiple 1GbE network connections to consolidated 10GbE.
- The QL41162HLRJ-CI can converge storage and networking I/O by deploying OS-based software iSCSI initiators over its 10GBASE-T connection.
- The Marvell FastLinQ QL41162HLRJ-CI Converged Network Adapter delivers a fully offloaded iSCSI and Fibre Channel over Ethernet (FCoE) solution that conserves CPU resources and delivers maximum performance.

Designed for Next-gen Server Virtualization

The QL41162HLRJ-CI supports today's most compelling set of powerful networking virtualization features: SR-IOV, NPAR, tunneling offloads (VXLAN, GRE, GENEVE, and NVGRE), and industry-leading performance, thus enhancing the underlying server virtualization features.

- SR-IOV delivers higher performance and lower CPU use with increased virtual machine (VM) scalability.
- Marvell NPAR enables up to 16 physical, switch-agnostic, switch-independent NIC partitions per adapter. Dynamic and fine-grained bandwidth provisioning enables control of network traffic from VMs and hypervisor services.
- Concurrent support for SR-IOV and NPAR enables virtual environments with the choice and flexibility to create an agile virtual server platform.
- Availability of both RSS and TSS allows for more efficient load balancing across multiple CPU cores.

High-Performance Multitenancy Delivered

As large-scale private and public cloud deployment requirements for isolation and security stretch the boundaries of traditional vLANs, the QL41162HLRJ-CI delivers network virtualization features for high-performance overlay networks.

- Designed to meet the demands of large, public cloud deployments, the QL41162HLRJ-CI Adapter features tunneling offloads for multitenancy with VXLAN, GRE, GENEVE, and NVGRE support.
- Line-rate 10GbE performance across individual ports in multitenant deployments maximizes server-processing performance by delivering an offloaded Ethernet adapter for enterprise, telco, and cloud deployments on Microsoft® Windows Server®, VMware vSphere®, and various Linux distributions.

Simplified Management

Marvell's QConvergeConsole (QCC) GUI delivers a broad set of powerful Ethernet and Fibre Channel (FC) adapter management features for administrators to maximize application performance and availability. QCC GUI offers application-based wizards to enable the environment to be quickly and easily provisioned based on published best practices. vCenter GUI and ESXCLI Plug-ins and OpenStack integration are also available.

QCS CLI is available for locally and remotely managing Linux and Windows servers. QCC PowerKit is available for remotely managing Linux, VMware (PowerCLI), and Windows servers. Additionally, pre-boot UEFI HII system BIOS device configuration is available on servers that support UEFI HII.

Accelerate Any Network With Universal RDMA Offload

The QL41162HLRJ-CI supports RoCE and iWARP acceleration to deliver low latency, low CPU utilization and high performance on Windows, VMware, and Linux operating systems. The QL41162HLRJ-CI Adapter has the unique capability to deliver Universal RDMA that enables RoCE, RoCEv2, and iWARP. Marvell Universal RDMA provides the ultimate flexibility in accelerating use cases like Microsoft Storage Spaces Direct (S2D), Windows Live Migration, Windows SMB Direct, Linux/Windows VF RDMA, VMware PVRDMA and vSAN, NVMe™ over Fabrics (NVMe-oF), CEPHS and NFS over RDMA, and so on. Marvell's cutting-edge offloading technology increases cluster efficiency and scalability to many thousands of nodes for HyperConverged infrastructure deployments. Customers looking to scale out NVMe-oF can leverage the QL41162HLRJ-CI capabilities of supporting NVMe-oF over TCP (NVMe/TCP) in addition to RDMA transports.

Accelerate Telco Network Function Virtualization (NFV) Workloads

In addition to OpenStack, the QL41162HLRJ-CI supports NFV, which allows decoupling network functions and services from dedicated hardware (such as routers, firewalls, and load balancers) into hosted VMs. NFV enables network administrators to flexibly create network functions and services as they need them, reducing capital expenditure and operating expenses, and enhancing business and network services agility. Marvell technology is integrated into the DPDK and can deliver up to 38 million packets per second to host the most demanding NFV workloads.

Trusted, Secure, Reliable, and Interoperable

The Marvell FastLinQ QL41162HLRJ-CI 10GbE Adapter adheres to standards that ensure interoperability with a wide range of network solutions. Marvell adapters are secure by design. Through public and private key encryption technology, the adapters enforce a process for secure firmware updates that prevent hackers from altering the code running on the adapters.

Host Bus Interface

Bus Interface

- PCI Express (PCIe) Gen 3 x8 (x8 physical connector)
- Supports PCIe upconfigure to reduce link width to conserve power

Host Interrupts

- MSI-X supports independent queues

I/O Virtualization and Multitenancy

- SR-IOV (up to 192 virtual functions)
- Switch-independent NPAR (up to 16 physical functions)
- GRE and NVGRE packet task offloads
- VXLAN packet task offloads
- GENEVE packet task offloads

Compliance

- *PCI Base Specification*, rev. 3.1
- *PCI Express Card Electromechanical Specification,* rev. 3.0
- *PCI Bus Power Management Interface Specification*, rev. 1.2
- *Advanced configuration and power interface (ACPI)* v2.0

Ethernet

Throughput

- 10Gbps line rate for dual port
- Auto negotiation: 1G/10G (BASE-T)

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

Ethernet (continued)

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)*
  - *(RJ45) 802.3-2018 Clauses 55 and 40 (10GBASE-T and 1000BASE-T)*
  - *1588-2002 PTPv1 (Precision Time Protocol)*
  - *1588-2008 PTPv2*
- RFCs
  - *IPv4 (RFC 791)*
  - *IPv6 (RFC 2460)*

Board Firmware Features

- Secure Firmware Update process

RDMA

Universal RDMA

- RoCE
- RoCEv2
- iWARP
- Storage over RDMA: iSER, SMB Direct, S2D, and NVMe-oF
- NFSoRDMA

RDMA Use Cases

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

FCoE-Offload

Performance

- 3.6 million FCoE IOPS

iSCSI-Offload

Performance

- 2.9 million iSCSI IOPS

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 and RESTful APIs for Linux, VMware, and Windows
- Pre-boot UEFI HII system BIOS device configuration pages
- Native OS management tools for networking

Boot Support

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

Operating System Support

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

Packaging

Ports

- Dual port

Form Factor

- MD2: PCI Express short, low-profile card: 167.65mm × 68.90mm (6.60in. × 2.71in.)

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

Environment and Equipment Specifications

Temperature

- Operating: 32°F to 131°F (0°C to 55°C)
- Storage: –40°F to 149°F (–40°C to 65°C)

Airflow

- Cooling requirements: 150/55 LFM/°C

Humidity (Relative, Non-condensing)

- Operating and non-operating: 10% to 90%

Compliance

- RoHS compliant

Cable Distance (Maximum)

Table 1. Cable Distance

Rate	Cable and Maximum Distance (m)
	RJ-45
10G	37 to 55 CAT6 and 100 CAT6a/7

RJ-45 = 10BASE-T

Approvals—Safety

US and Canada

- UL 60950-1
- CSA C22.2.

Europe

- TUV EN60950-1
- TUV IEC 60950-1
- CB Certified

Agency Approvals—EMI and EMC

US and Canada

- FCC Rules, CFR Title 47, Part 15, Subpart Class A
- Industry Canada, ICES-003: Class A

Europe

- EN55032
- EN55024
- EN61000-3-2
- EN61000-3-3

Japan

- VCCI: Class A

New Zealand and Australia

- AS/NZS: Class A

Korea

- KC-RRA Class A

Taiwan

- BSMI CNS 13438

Cisco Ordering Information

UCSC-PCIE-QD10GC,
UCSC-PCIE-QD10GC= (spare model number)

- Dual Port 10GbE RJ-45 Adapter
- Form factor: PCIe, MD2
- Platform: L2, RoCE/RocEv2, iWARP, iSCSI, FCoE
- All adapters support adaptive voltage scaling (AVS).
- Compatible with Cisco UCS C4200 Series Servers.
- Adapters are purchased as a bulk kit (BK).
- Adapters ship with:
  - Standard-size bracket installed. A low-profile bracket is also included
  - RJ-45 connectors. Intended for use with twisted pair copper cabling (not included)