


# Marvell® FastLinQ® QL45611HLCU Adapter

Multiport 10/25/100Gb Ethernet Adapter with Universal RDMA



**QL45611HLCU**

- Single port 100GbE Adapter with quad 10/25GbE connectivity using breakout cables
- Delivers full line-rate performance across a single port
- 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

The Marvell FastLinQ QL45611HLCU Intelligent Ethernet Adapter with Universal Remote Direct Memory Access (RDMA) supports 10-Gigabit Ethernet (GbE), 25GbE, and 100GbE QSFP28, as well as LAN (TCP/IP) traffic at line-rate speeds. The QL45611HLCU provides extremely low host CPU usage by enabling full stateless offloads to meet the performance requirements of the most demanding enterprise applications.

The Marvell FastLinQ QL45611HLCU 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 10/25/100GbE bandwidth, the QL45611HLCU Intelligent Ethernet Adapter offers switch-independent NIC partitioning (NPAR), which enables segmentation of a 10/25/100GbE 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 10/25/100GbE solution that boosts CPU efficiency and reduces capital expenditures (CAPEX) and operational expenditures (OPEX) of the migration to 10/25/100GbE. The Marvell FastLinQ QL45611HLCU is the best choice for workload-intensive computing environments, providing reliable, high-performance 10/25/100GbE connectivity solutions.

## Features

- PCI Express® (PCIe®) Gen 3 x16 (8GT/s) support
- Full line-rate performance across single and quad ports
- Broad operating system (OS) and hypervisor support
- Network boot support
  - iSCSI (SW) 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 68 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)
- MSI and MSI-X support
- IPv4 and IPv6 stateless offloads
- PCI-SIG® single root input/output virtualization (SR-IOV) with up to 240 virtual functions
  - Linux SR-IOV VF QoS support through the **ip link set** command
- Comprehensive stateless offloads
- Auto negotiation: 10G/25G/100G
- FastLinQ SmartAN™ for simplified connectivity with SFP and QSFP interfaced switches. (SFP/QSFP interfaces can accept direct attach cable (DAC) or optical (discrete separate optical modules or active optic cable (AOC)) connections).
- RX/TX multiqueue
  - VMware® NetQueue
  - 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)
  - 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
- Offloaded storage over Ethernet
  - Increases server performance with full hardware offload for storage traffic

## Benefits

### Simplified Migration to 10/25/100GbE

The Marvell FastLinQ QL45611HLCU 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 10/25/100GbE 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, QL45611HLCU Adapters can further partition their network bandwidth into multiple virtual connections, making 1 adapter appear as either 8 (in 1x100G mode) or 16 (in quad port mode) 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 10/25/100GbE migration.
- The QL45611HLCU is the ideal choice for migrating multiple 1GbE network connections to consolidated 10/25/100GbE. It is compatible with SR optics, LR optics, direct-attach copper (DAC) at 10 and 100G; as well as SR optics, LR optics, and DAC that does not require RS-FEC at 25G.
- The QL45611HLCU can converge storage and networking I/O by deploying OS-based software iSCSI initiators over its optical or DAC connections.

### Designed for Next-gen Server Virtualization

The Marvell FastLinQ QL45611HLCU 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 8 (1x100G mode) or 16 (quad port mode) 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 Marvell FastLinQ QL45611HLCU delivers network virtualization features for high-performance overlay networks.

- Designed to meet the demands of large, public cloud deployments, the QL45611HLCU features tunneling offloads for multitenancy with VXLAN, GRE, GENEVE, and NVGRE support.
- Line-rate 10/25/100GbE 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) delivers a broad set of powerful Ethernet adapter management features for administrators to maximize application performance and availability. vCenter GUI and ESXCLI Plug-ins 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 vCenter (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 Marvell FastLinQ QL45611HLCU supports RoCE and iWARP (capable, not enabled) acceleration to deliver low latency, low CPU utilization and high performance on Windows, VMware, and Linux operating systems. The QL45611HLCU 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, VMware PVRDMA and vSAN, NVMe™ over Fabrics (NVMe-oF), CEPHS and NFS over RDMA, Linux/Windows VF 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 QL45611HLCU capabilities of supporting RoCE, RoCEv2, iWARP, or the newly added TCP (referred to as *NVMe-oF over TCP* or *NVMe/TCP*) in addition to the RDMA transport fabrics.

## Accelerate Telco Network Function Virtualization (NFV) Workloads

The Marvell FastLinQ QL45611HLCU 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 68 million packets per second to host the most demanding NFV workloads.

The QL45611HLCU also supports the NSX-T/N-VDS Enhanced data path/Network Stack (ENS) polling mode driver (QDeNTV\_ens) for NFV workloads on VMware ESXi 6.7.

### **OPEX Savings With Low-power PCIe Gen 3**

The QL45611HLCU Adapter is a PCIe Gen 3 host bus connectivity based adapter that has one of the lowest power-consumption profiles in the industry.

- Supporting PCIe Gen 3 host bus connectivity enables the QL45611HLCU to deliver line-rate performance without compromise.
- The QL45611HLCU is designed to provide maximum power efficiency while delivering a high-performance, I/O connectivity platform.

### **Trusted, Secure, Reliable, and Interoperable**

The Marvell FastLinQ QL45611HLCU 10/25/100GbE 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 x16

### Host Interrupts

- MSI-X supports independent queues

### I/O Virtualization and Multitenancy

- SR-IOV (up to 240 virtual functions)
- Switch-independent NPAR (up to 8 physical functions in 1x100G mode or 16 physical functions in quad port mode)
- 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

- 10/25/100Gbps line rate for single and quad port
- 10G/25G/100G 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 hardware 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)/RSS version 2 (RSSv2). Linux Multiqueue, and Virtual Machine Device queues (VMDq)
- DPDK/VMware N-VDS Enhanced Data Path
- 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)
  - (25GbE) 802.3-2018 Clause 110 (Direct Attach Copper)
  - (10GbE) 802.3-2018 Clause 52 (10Gb Ethernet Optical), Annex 83B (XLAUI Chip to module) to support 802.3-2018 Clause 52 (SR and LR optical)
  - (100GbE) 802.3-2018 Clause 92 (Direct Attach Copper), Annex 83E (chip to module for use with Clause 95 (SR optical), Clause 88 (LR optical), and 100G PSM4 and CWDM MSA modules)
  - 1588-2002 PTPv1 (Precision Time Protocol)
  - 1588-2008 PTPv2
  - (10GbE) 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)

## RDMA

### Universal RDMA

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

### RDMA Use Cases

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

## 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

## 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

### Boot Support

- PXE 2.0
- UEFI
- iSCSI (SW) remote boot

### Operating System Support

- For the latest applicable operating system information, see [www.marvell.com](http://www.marvell.com)

## Packaging

### Ports

- Single and quad port modes available over a single QSFP port. See the adapter features on page 8.

### Form Factor

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

## Environment and Equipment

### Temperature

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

### Airflow

- 200LFM

### Humidity (Relative, Non-condensing)

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

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 (continued)**

**Compliance**

- RoHS compliant

**Cable Distance (Maximum)**

**Table 1. Cable Distance**

Rate	Cable and Maximum Distance (m)		
	DAC	SR FOC	AOC
10G	7	400 OM4 300 OM3	30
25G	3 <sup>1</sup>	Note 1	Note 1
100G <sup>2</sup>	5	100 OM4 70 OM3	30

DAC = Direct attach cable  
SR FOC = SR fiber optic cable  
AOC = Active optic cable

1. Only for modules and cables that do not require RS-FEC. The actual maximum supported lengths on 25G DACs and optics depend on the module vendor's specifications. For DACs, the gauge and cable quality (CA-25G-N versus CA-25G-S) affect the usable maximum distance when no FEC or FC-FEC is used: CA-25G-N cables support 3m length with no FEC, CA-25G-S cables support 3m length with FC-FEC. For AOCs, 30m is generally supported with no FEC or FC-FEC. For separate 25G SR optical modules, the vendor specifications and the type of FEC (no FEC versus FC-FEC) and the type of fiber (OM3 versus OM4) determines the maximum reach possible: generally 30m on OM3/60m on OM4 with no FEC, and 70m on OM3/100m on OM4 with FC-FEC. LR optical modules may need FEC to reach their maximum 10KM distance.

2. Maximum distances on 100G DACs and SR/LR optics may require RS-FEC.

**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-RRR Class A

**Taiwan**

- BSMI CNS 13438

**Table 2. QL45611HLCU Features**

General Specs	
Ports	1/4
Port Speeds	1x100, 4x25 <sup>1</sup> , 4x10 <sup>1</sup>
Connectors	QSFP28
Form Factor (see “Packaging” on Page 6)	PCIe MD2
Media	DAC <sup>2</sup> , optics <sup>2</sup> , AOC <sup>2</sup>
FEC Mode	FireCode and Reed Solomon FEC <sup>3</sup>
SmartAN™ Mode	√
Storage	
Universal RDMA (RoCE/RoCEv2/ iWARP)	√ (Note 4)
NVME-oF over TCP/RDMA	√
FCoE Offload	—
iSCSI Offload	—
Virtualization and Cloud	
Concurrent SR-IOV/NPAR	√ (Note 5)
DPDK	√
Flow Filtering	√
Tunneling Offload (VXLAN/GENEVE/ NVGRE/GRE)	√
Physical Specifications	
Cooling Requirements (LFM/°C)	N/A (Note 6)

1. Quad port mode over a single QSFP interface.
2. Long DACs and SR optics/AOCs that require Reed Solomon FEC on 25GbE links are not supported.
3. Reed Solomon FEC supported on 100GbE links only.
4. 1x100G mode is iWARP capable. Contact your Marvell representative for more details. The QL45611HLCU supports iWARP (and RoCE/RoCEv2) in 4x10G and 4x25G modes.
5. There are up to eight NPAR partitions per port in single port mode, and up to four NPAR partitions per port in quad-port mode.
6. Heatsinks have a built-in cooling fan.

**Table 3. Ordering Information**

Part Number	Description	Form Factor	Platform
QL45611HLCU	Single Port 100GbE/Quad Port 10GbE/Quad Port 25GbE QSFP28 Adapter	PCIe MD2	L2, RoCE/RoCEv2, iWARP capable

This adapter supports adaptive voltage scaling (AVS).

To order a bulk kit, add *-BK* to the end of the part number; for example, QL45611HLCU-BK.

To order a single pack, add *-CK* to the end of the part number; for example, QL45611HLCU-SP.

All *HL* (MD2 PCIe stand-up) adapters come with both full height and low profile brackets.

DAC cables, SR/LR optics are not included. See <https://www.marvell.com/documents/gx4rrv7xnip1ea6irzc7/> for a list of cables and optics that have been tested by Marvell and its partners.





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.

Copyright © 2021 Marvell. All rights reserved. Marvell and the Marvell logo are trademarks of Marvell or its affiliates. Please visit [www.marvell.com](http://www.marvell.com) for a complete list of Marvell trademarks. Other names and brands may be claimed as the property of others.