


# MARVELL FastLinQ 45000 Series

Multiport 10/25/40/50/100Gb Ethernet Adapters with Universal RDMA



- Single, dual, and quad port 10/25/40/50/100GbE Adapter options
- Delivers full line-rate performance across single, dual, and quad ports
- 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
- Supports iSCSI-Offload and FCoE-Offload

The Marvell® FastLinQ® 45000 Series Intelligent Ethernet Adapters with Universal Remote Direct Memory Access (RDMA)—available in 10-Gigabit Ethernet (GbE) SFP+, 25GbE SFP28, 40GbE QSFP+, 50GbE QSFP28, and 100GbE QSFP28—support LAN (TCP/IP) traffic at line-rate speeds. The 45000 Series 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 45000 Series 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 10/25/40/50/100GbE bandwidth, the 45000 Series Intelligent Ethernet Adapter offers switch-independent NIC partitioning (NPAR), which enables segmentation of a single 10/25/40/50/100GbE 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 10/25/40/50/100GbE solution that boosts CPU efficiency and reduces capital expenditures (CAPEX) and operational expenditures (OPEX) of the migration to 10/25/40/50/100GbE. The Marvell FastLinQ 45000 Series Adapters are the best choice for workload-intensive computing environments, providing reliable, high-performance 10/25/40/50/100GbE connectivity solutions.

## FEATURES

- PCI Express® (PCIe®) Gen 3 x16 (8GT/s) support
- Full line-rate performance across single, dual, and quad 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 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
- Comprehensive stateless offloads
- Auto negotiation: 1G/10G/25G/40G/50G/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
  - Industry-leading FCoE-Offload performance of up to 7 million IOPS, suitable for high-density server virtualization and large databases
  - Industry-leading iSCSI-Offload performance of up to 5.4 million IOPS, suitable for a diverse set of applications leveraging the flexibility of iSCSI

## BENEFITS

### Simplified Migration to 10/25/40/50/100GbE

Marvell FastLinQ 45000 Series Adapters feature 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 10/25/40/50/100GbE 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, 45000 Series Adapters can further partition their 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 10/25/40/50/100GbE migration.
- Available in SR optics, LR optics, and direct-attach copper (DAC), 45000 Series Adapters are the ideal choice for migrating multiple 1GbE network connections to consolidated 10/25/40/50/100GbE.
- Marvell FastLinQ 45000 Series Adapters can converge storage and networking I/O by deploying OS-based software iSCSI initiators over their optical or DAC connections.
- Marvell 45000 Series Converged Network Adapters (CNAs) deliver 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 Marvell FastLinQ 45000 Series Adapters support 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 Marvell FastLinQ 45000 Series Adapters deliver network virtualization features for high-performance overlay networks.

- Designed to meet the demands of large, public cloud deployments, the 45000 Series Adapters feature tunneling offloads for multitenancy with VXLAN, GRE, GENEVE, and NVGRE support.
- Line-rate 10/25/40/50/100GbE 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) 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 Marvell FastLinQ 45000 Series Adapters support RoCE and iWARP acceleration to deliver low latency, low CPU utilization and high performance on Windows, VMware, and Linux operating systems. The 45000 Series Adapters have 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 45000 Series 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

In addition to OpenStack, the Marvell FastLinQ 45000 Series Adapters support 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.

### OPEX Savings with Low-power PCIe Gen 3

The 45000 Series are PCIe Gen 3 based adapters that have one of the lowest power-consumption profiles in the industry.

- Supporting the latest generation of host bus connectivity, PCIe Gen 3 enables the Marvell FastLinQ 45000 Series Adapters to deliver line-rate, dual-port performance without compromise.
- 45000 Series Adapters are designed to provide maximum power efficiency while delivering a high-performance, I/O connectivity platform.

### TRUSTED, SECURE, RELIABLE, AND INTEROPERABLE

The FastLinQ 45000 Series 10/25/40/50/100GbE Adapters adhere 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 (QL45600 or QL45400) or Gen 3 x8 (QL45200) Series

### Host Interrupts

- MSI-X supports independent queues

### I/O Virtualization and Multitenancy

- SR-IOV (up to 240 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
- Open Compute Project, OCP Mezzanine card 2.0 Design Specification, v1.00

## Ethernet

### Throughput

- 10/25/40/100Gbps line rate for single, dual, and quad port
- 1G/10G/25G/40G/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 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)

### Stateless Offload (continued)

- 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)
  - (25GbE SFP28) 802.3-2018 Clause 110 (Direct Attach Copper)
  - (10GbE SFP+) 802.3-2018 Clause 52 (10Gb Ethernet Optical)
  - (40GbE QSFP+) 802.3-2018 Clause 85 (Direct Attach Copper), Clause 86 (SR optical), Clause 87 (LR optical)
  - (100GbE QSFP28) 802.3-2018 Clause 92 (Direct Attach Copper), Clause 95 (SR optical), Clause 88 (LR optical), Annex 83E (chip to module for use with 100G PSM4 and CWDM MSA modules)
  - 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)

## 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

### RDMA Use Cases (continued)

- SMB Direct
- NVMe-oF
- NFS
- CEPHS over RDMA
- VF RDMA

## FCoE-Offload

### Performance

- 7 million FCoE IOPS

## iSCSI-Offload

### Performance

- 5.4 million iSCSI IOPS

## Forward Error Correction (FEC)

- FireCode, BASE-R IEEE 802.3-2018 Clause 74, or FC-FEC
- Reed Solomon IEEE 802.3-2018 Clause 91 (for 100G only)

## 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 remote boot
- FCoE boot from SAN

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

**Operating System Support**

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

**Packaging****Ports**

- Single, dual, and quad port variants available. See the list of adapters and their features on page 7.

**Form Factor**

- MD2: PCI Express short, low-profile card: 167.65mm × 68.90mm (6.60in. × 2.71in.)
- OCP 2.0: Complies with Open Compute Project (OCP) 2.0

**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**

- See the table on page 7.

**Humidity (Relative, Non-condensing)**

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

**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
40G	7	150 OM4 100 OM3	30
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 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: 1m 25G DACs usually do not require FEC; 2m 25G DACs may require FC-FEC; 3m and beyond 25G DACs may require RS-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.

**Compliance**

- RoHS compliant

**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



**Table 2. Features**

Adapter Name	QL45211 HLCU	QL45212 HLCU	QL45411 HLCU	QL45412 HLCU	QL45462 HLCU	QL45611 HLCU	QL45631 HOCU
<b>General Specs</b>							
Ports	1	2	1	2	2	1	1
Port Speeds	1, 10, 25	1, 10, 25	1x40, 4x10 <sup>1</sup>	2x40, 4x10 <sup>1</sup>	2x40, 4x10 <sup>1</sup>	1x100, 4x25 <sup>1</sup> , 4x10 <sup>1</sup>	1x100, 4x25 <sup>1</sup> , 4x10 <sup>1</sup>
Connectors	SFP, SFP+, SFP28	SFP, SFP+, SFP28	QSFP+	QSFP+	QSFP+	QSFP28	QSFP28
Form Factor	PCIe MD2	PCIe MD2	PCIe MD2	PCIe MD2	PCIe MD2	PCIe MD2	OCF 2.0
Media	DAC <sup>2</sup> , optics <sup>2</sup> , AOC <sup>2</sup>	DAC <sup>2</sup> , optics <sup>2</sup> , AOC <sup>2</sup>	DAC, optics, AOC	DAC, optics, AOC	DAC, optics, AOC	DAC <sup>2</sup> , optics <sup>2</sup> , AOC <sup>2</sup>	DAC <sup>2</sup> , optics <sup>2</sup> , AOC <sup>2</sup>
FEC Mode	Firecode FEC	Firecode FEC	Firecode FEC	Firecode FEC	Firecode FEC	Firecode and Reed Solomon FEC <sup>3</sup>	Firecode and Reed Solomon FEC <sup>3</sup>
SmartAN™ Mode	✓	✓	✓	✓	✓	✓	✓
<b>Storage</b>							
Universal RDMA (RoCE/ RoCEv2/ iWARP)	✓	✓	✓	✓	✓	✓ <sup>4</sup>	✓ <sup>4</sup>
NVME-oF over TCP/ RDMA	✓	✓	✓	✓	✓	✓	✓
FCoE Offload	—	—	—	—	✓	—	—
iSCSI Offload	—	—	—	—	✓	—	—
<b>Virtualization and Cloud</b>							
Concurrent SR-IOV/ NPAR <sup>5</sup>	✓	✓	✓	✓	✓	✓	✓
DPDK	✓	✓	✓	✓	✓	✓	✓
Flow Filtering	✓	✓	✓	✓	✓	✓	✓
Tunneling Offload (VXLAN/GENEVE/ NVGRE/GRE)	✓	✓	✓	✓	✓	✓	✓
<b>Physical Specifications</b>							
Cooling Requirements (LFM/°C)	N/A <sup>6</sup>	N/A <sup>6</sup>	N/A <sup>6</sup>	N/A <sup>6</sup>	N/A <sup>6</sup>	N/A <sup>6</sup>	350/45

1. Quad port mode over a single QSFP interface. On dual-port QSFP adapters, the second QSFP port is disabled.
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.
5. There are up to eight NPAR partitions per port in single- and dual-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
QL45631HOCU	Single Port 100GbE/Quad Port 10GbE/ Quad Port 25GbE QSFP28 Adapter	OCP 2.0	L2, RoCE/RocEv2, iWARP capable
QL45611HLCU	Single Port 100GbE/Quad Port 10GbE/ Quad Port 25GbE QSFP28 Adapter	PCIe MD2	L2, RoCE/RocEv2, iWARP capable
QL45462HLCU	Dual Port 40GbE/Quad Port 10GbE QSFP+ Adapter	PCIe MD2	L2, RoCE/RocEv2, iWARP, iSCSI, FCoE
QL45412HLCU	Dual Port 40GbE/Quad Port 10GbE QSFP+ Adapter	PCIe MD2	L2, RoCE/RocEv2, iWARP
QL45411HLCU	Single Port 40GbE/Quad Port 10GbE QSFP+ Adapter	PCIe MD2	L2, RoCE/RocEv2, iWARP
QL45212HLCU	Dual Port 25/10GbE SFP28 Adapter	PCIe MD2	L2, RoCE/RocEv2, iWARP
QL45211HLCU	Single Port 25/10GbE SFP28 Adapter	PCIe MD2	L2, RoCE/RocEv2, iWARP

All adapters support 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, QL45462HLCU-SP.

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

SFP/QSFP, DAC cables, SR/LR optics are not included. See <https://www.marvell.com/documents/gx4rrv7xniplea6irzc7/> for a list of cables and optics that have been tested by Marvell and its partners.



Marvell\_45000 Series

**ABOUT MARVELL:** 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, processing, networking, security 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](http://www.marvell.com).

Copyright © 2019 Marvell. All rights reserved. Marvell, the Marvell logo, FastLinQ, QLogic, QConvergeConsole and SmartAN, are registered trademarks or trademarks of Marvell. All other trademarks are the property of their respective owners.