


Marvell® FastLinQ® 41000 Series

Multiport 10/25Gb Ethernet Adapters with Universal RDMA



- Single, dual, and quad port 10GbE and 25GbE Adapter options
- Delivers full line-rate 10/25GbE performance across single and dual 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
- 10GBASE-T version provides low-cost and easy-to-install RJ45 connectivity that is compatible with existing 1GbE

The FastLinQ 41000 Series Intelligent Ethernet Adapter with Universal Remote Direct Memory Access (RDMA)—available in 10GBASE-T (RJ45), 10-Gigabit Ethernet (GbE) SFP+, and 25GbE SFP28—supports LAN (TCP/IP) traffic at 10/25GbE line-rate speeds. The adapters deliver true 10GbE and 25GbE speed, power, and performance. Integrated, advanced networking eliminates I/O bottlenecks and conserves CPU cycles.

The FastLinQ 41000 Series leverages Marvell's 15+ years of expertise in Ethernet by providing the highest levels of performance, efficiency, and scalability for Open Compute server and storage applications in Web 2.0, enterprise data centers, and cloud infrastructure.

For more effective use of the 10/25GbE bandwidth, the 41000 Series Intelligent Ethernet Adapter offers switch-independent NIC partitioning (NPAR), which enables segmentation of a single 10/25GbE 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 10/25GbE solution that boosts CPU efficiency and reduces capital expenditures (CAPEX) and operational expenditures (OPEX) of the migration to 10/25GbE. The FastLinQ 41000 Series Adapters are the best choice for workload-intensive computing environments, providing reliable, high-performance 10/25GbE connectivity solutions.

FEATURES

- PCI Express® (PCIe®) Gen 3 x8 (8GT/s) support
- Full line-rate performance across single and dual 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 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 (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 (BASE-T) and 10G/25G (on direct attach cable (DAC) cable using 10GBASE-KR/25GBASE-CR)
- FastLinQ SmartAN™ for simplified connectivity with 10G SFP+/25G SFP28 interfaced switches. (SFP+ interfaces can accept DAC or optical (discrete 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)
- RSSv2
- 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 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

BENEFITS

Simplified Migration to 10/25GbE

FastLinQ 41000 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 10/25GbE 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, 41000 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 10/25GbE migration.
- Available in 10GBASE-T, SR and LR optics, and direct-attach copper (DAC), 41000 Series Adapters are the ideal choice for migrating multiple 1GbE network connections to consolidated 10/25GbE.
- FastLinQ 41000 Series Adapters can converge storage and networking I/O by deploying OS-based software iSCSI initiators over their 10GBASE-T and optical or DAC connections.
- Marvell 41000 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 FastLinQ 41000 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 FastLinQ 41000 Series Adapters deliver network virtualization features for high-performance overlay networks.

- Designed to meet the demands of large, public cloud deployments, the 41000 Series Adapters feature tunneling offloads for multitenancy with VXLAN, GRE, GENEVE, and NVGRE support.
- Line-rate 10/25GbE 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 FastLinQ 41000 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 41000 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, 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 41000 Series 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 FastLinQ 41000 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 38 million packets per second to host the most demanding NFV workloads.

Trusted, Secure, Reliable, and Interoperable

The FastLinQ 41000 Series 10/25GbE 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 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
- OCP NIC 3.0 Design Specification (draft) v1.0

Ethernet

Throughput

- 10/25Gbps line rate for single and dual port
- Auto negotiation: 1G/10G (BASE-T) and 10G/25G (on DAC cable using 10GBASE-KR/25GBASE-CR)

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)
 - (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)
 - (RJ45) 802.3-2018 Clauses 55 and 40 (10GBASE-T and 1000BASE-T)
 - 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, 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

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 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 boot from SAN

Operating System Support

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

Packaging

Ports

- Single, dual, and quad port variants available. See the list of adapters and their features in [Table 2](#).

Form Factor

- MD2: PCI Express short, low-profile card: 167.65mm × 68.90mm (6.60in. × 2.71in.)
- OCP 3.0 Complies with Open Compute Project (OCP) 3.0. Adapter equipped with pull tabs.

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

- See [Table 2](#).

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)			
	DAC	SR FOC	AOC	RJ-45
10G	7	400 OM4 300 OM3	30	37 to 55 CAT6 100 CAT6a/75
25G	5	100 OM4 70 OM3	30	—

DAC = Direct attach cable
SR FOC = SR fiber optic cable
AOC = Active optic cable
RJ-45 = 10BASE-T variants only

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

UK

- UKCA
- BS DoC

Table 2. Features

Adapter Name	QL41132 HLRJ	QL41134 HLRJ	QL41132 HQRJ	QL41134 HQRJ	QL41132 HLCU	QL41134 HLCU	QL41254 HLCU	QL41234 HLCU	QL41212 HLCU	QL41232 HQCU	QL41134 HQCU	QL41262 HLCU
General Specs												
Ports	2	4	2	4	2	4	4	4	2	2	4	2
Port Speeds	1, 10	1, 10	1, 10	1, 10	1 ¹ , 10	1 ¹ , 10	1 ¹ , 10, 25	1 ¹ , 10,25	1 ¹ , 10,25	1 ¹ , 10,25	1 ¹ , 10	1 ¹ , 10,25
Connectors	BASE-T ²	BASE-T ²	BASE-T ²	BASE-T ²	SFP ¹ , SFP+	SFP ¹ , SFP+	SFP ¹ , SFP+, SFP28	SFP ¹ , SFP+, SFP28	SFP ¹ , SFP+, SFP28	SFP ¹ , SFP+, SFP28	SFP ¹ , SFP+	SFP ¹ , SFP+, SFP28
Form Factor	PCIe MD2	PCIe MD2	OCP 3.0	OCP 3.0	PCIe MD2	PCIe MD2 ³	PCIe MD2 ³	PCIe MD2 ³	PCIe MD2	OCP 3.0	OCP 3.0	PCIe MD2
Media	RJ-45 CAT	RJ-45 CAT	RJ-45 CAT	RJ-45 CAT	DAC, optics, AOC	DAC, optics, AOC	DAC, optics, AOC	DAC, optics, AOC	DAC, optics, AOC	DAC, optics, AOC	DAC, optics, AOC	DAC, optics, AOC
IEEE 802.3az (EEE)	√	√	√	√	—	—	—	—	—	—	—	—
SmartAN™ Mode	—	—	—	—	√	√	√	√	√	√	√	√
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	√	√	√	√	√	√	√	√	√	√	√	√
Tunneling Offload (VXLAN/GENEVE/NVGRE/GRE)	√	√	√	√	√	√	√	√	√	√	√	√
Physical Specifications												
Cooling Requirements (LFM/°C)	150/55	250/55	225/55	275/55	100/55	150/55	150/55	150/55	100/55	200/55	250/55	100/55

1. Available in a future release.

2. BASE-T (RJ-45) interfaces support Auto-Negotiation, 1GbE full duplex, and 10GbE full duplex.

3. Compatible with PCIe MD2 form factor, but not fully compliant.




Table 3. Ordering Information

Part Number	Description	Form Factor	Platform
QL41234HLCU-BK/SP/CK	Quad Port 25/10GbE SFP28 Adapter	PCIe MD2 ¹	L2, RoCE/RoCEv2, iWARP
QL41254HLCU-BK/SP/CK	Quad Port 25/10GbE SFP28 Adapter	PCIe MD2 ¹	L2, RoCE/RoCEv2, iWARP, iSCSI
QL41134HLCU-BK/SP/CK	Quad Port 10GbE SFP+ Adapter	PCIe MD2 ¹	L2, RoCE/RoCEv2, iWARP
QL41134HLRJ-BK/SP/CK	Quad Port 10GbE RJ-45 Adapter	PCIe MD2	L2, RoCE/RoCEv2, iWARP
QL41212HLCU-11-BK/SP/CK	Dual Port 25/10GbE SFP28 Adapter	PCIe MD2	L2, RoCE/RoCEv2, iWARP
QL41262HLCU-11-BK/SP/CK	Dual Port 25/10GbE SFP28 Adapter	PCIe MD2	L2, RoCE/RoCEv2, iWARP, iSCSI, FCoE
QL41132HLRJ-11-BK/SP/CK	Dual Port 10GbE RJ-45 Adapter	PCIe MD2	L2, RoCE/RoCEv2, iWARP
QL41132HLCU-11-BK/SP/CK	Dual Port 10GbE SFP+ Adapter	PCIe MD2	L2, RoCE/RoCEv2, iWARP
QL41134HQCU-BK/SP/CK	Quad Port 10GbE SFP+ Adapter	OCP 3.0 ²	L2, RoCE/RoCEv2, iWARP
QL41134HQRJ-BK/SP/CK	Quad Port 10GbE RJ-45 Adapter	OCP 3.0 ²	L2, RoCE/RoCEv2, iWARP
QL41132HQRJ-BK/SP	Dual Port 10G RJ-45 Adapter	OCP 3.0 ²	L2, RoCE/RoCEv2, iWARP
QL41232HQCU-BK/SP	Dual Port 25/10GbE SFP28 Adapter	OCP 3.0 ²	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, QL41262HLCU-11-BK.

To order a single pack, add *-SP* to the end of the part number; for example, QL41262HLCU-11-SP.

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

Twisted pair cabling, DAC cables, SR/LR optics are not included. See <https://www.marvell.com/documents/xalfardzafh32cfvi0z/> for a list of cables and optics that have been tested by Marvell and its partners.

10BGASE-T variants ship with RJ45 connectors. Intended for use with twisted pair copper cabling (not included).

1. Compatible with PCIe MD2 form factor, but not fully compliant.

2. OCP adapters have a pull tab thumbscrew.



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.

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