


Marvell® FastLinQ® 41000 Series Adapters

Broad Portfolio of 10GbE and 25GbE Ethernet Adapters with Universal RDMA for HPE® ProLiant and Apollo Gen10 Plus Servers



- Dual and quad port 10GbE and dual port 25GbE adapter options
- Delivers full line-rate 10/25GbE performance across all 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 server virtualization technologies such as SR-IOV
- Boosts host CPU efficiency with hardware offloads for GRE, NVGRE, GENEVE, and VXLAN tunnels
- 10GBASE-T version provides low-cost and easy-to-install RJ-45 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 (RJ-45), 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. Optimized for use with HPE ProLiant® Gen10 Plus Servers across enterprises, managed service providers (MSPs), and large public and scalable public cloud deployments, the 41000 Series enables organizations to achieve new levels of performance in physical, virtual, and cloud environments.

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 Marvell FastLinQ 41000 Series Adapters are the best choice for workload-intensive computing environments, providing reliable, high-performance 10/25GbE connectivity solutions.

Marvell features that collectively deliver these state-of-the-art network adapters include:

- Cutting-edge server virtualization technology—single-root I/O virtualization (SR-IOV)
- Network virtualization—offloads for Virtual Extensible LAN (VXLAN), Generic Network Virtualization Encapsulation (GENEVE), Generic Routing Encapsulation (GRE), and Network Virtualization using Generic Routing Encapsulation (NVGRE)
- Universal RDMA technologies—RDMA over Converged Ethernet (RoCE), RoCEv2, and Internet wide area RDMA protocol (iWARP)
- NVMe over Fabric (NVMe-oF™) support for NVMe over RoCE/RoCEv2, NVMe over iWARP, and NVMe over TCP
- Extremely low host CPU usage by enabling full hardware offloads
- Secure firmware update with private/public key encryption to prevent rogue firmware installations and enhance security

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 (software) remote boot with software initiators (no iSCSI hardware offload)
 - 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, as well as QCC vCenter GUI and ESXCLI Plug-ins
- Switch-independent NPAR with up to 8 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—RoCE, RoCEv2, and 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® 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 direct attach cable (DAC) or optical (discrete or active optic cable (AOC)) connections.
- RX/TX multiqueue
 - VMware® NetQueue
 - Windows® Hyper-V® Dynamic Virtual Machine Queue
 - Linux Multiqueue
- Tunneling offloads
 - Windows NVGRE
 - Linux GRE
 - VMware, Windows, and Linux VXLAN
 - Linux and VMware GENEVE
- Receive side scaling (RSS)
- RSSv2
- Virtual switch RSS (vRSS)
- Dynamic Virtual Machine Multiqueues (VMMQ)
- 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 Data 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)
- Data Center Quantized Congestion Notification (DCQCN)
- Non-offloaded Storage over Ethernet
 - iSCSI using OS-based software initiators

Benefits

Reduce Capital Expenditure and Operating Expense

FastLinQ 41000 Series Adapters enable cloud providers and large-scale data center operators to reduce operating expense while continuing to scale their network of server and storage nodes to meet increasing demands. Marvell 10GBASE-T technology is cost-efficient and power-efficient, supporting Energy Efficient Ethernet™ (EEE) to reduce idle power consumption.

Simplified Migration to 10/25GbE

The 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 8 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 DAC, 41000 Series Adapters are the ideal choice for migrating multiple 1GbE network connections to consolidated 10/25GbE.
- FastLinQ 41000 Series Adapters support SmartAN technology that automatically configures the adapter when connections change from 10GbE to 25GbE.
- Simplifies transition to new networks by supporting 10GbE and 25GbE speeds independently on each port of 10/25GbE adapters.
- 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.

Designed for Next-gen Server Virtualization

The Marvell 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 8 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 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.

Marvell 41000 Series Adapters also support Redfish Device Enablement (RDE) for integration into HPE-specific management utilities like HME Integrated Lights Out (iLO) remote server management.

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)—Azure Stack Hyper Converged Infrastructure (HCI) and SMB Direct, Windows Live Migration, VMware PVRDMA and vSAN, NVMe™ over Fabrics (NVMe-oF), CEPHS and NFS over RDMA, and so on. Marvell's offloading technology increases cluster efficiency and scalability for HyperConverged infrastructure deployments. The 41000 Series Adapters support NVMe/RoCE—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

The 41000 Series Adapters support NFV, which allows decoupling of 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. The 41000 Series Adapters are integrated into the DPDK and can deliver up to 38 million packets 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 8 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 3.0 Design Specification*, v1.0

Ethernet

Throughput

- 10/25Gbps line rate across all ports
- Dual and quad port 10Gbps and dual port 25Gbps
- 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)

Ethernet (continued)

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)/VMMQ/vRSS/RSSv2, 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)*
 - *(RJ-45) 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)*
 - *(10GbE SFP+) 802.3-2018 Clause 52 (10Gb Ethernet Optical)*
 - *(RJ-45) 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 (Azure Stack HCI), and NVMe-oF
- NFSoRDMA

RDMA Use Cases

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

Forward Error Correction (FEC)

- FireCode "Base-R" IEEE802.3-2018 Clause 74 (FC-FEC)
- Reed Solomon IEEE 802.3-2018 Clause 91 (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 for Linux, VMware, and Windows
- Pre-boot UEFI HII system BIOS device configuration pages
- Native OS management tools for networking
- Redfish Device Enablement (RDE), read only

Boot Support

- PXE 2.0
- UEFI
- iSCSI remote boot

Operating System Support

- For the latest applicable operating system information, see <https://www.hpe.com/us/en/servers/server-operating-systems.html>

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

Packaging

Ports

- Dual and quad port variants available. See the list of adapters and their features in [Table 2](#) and [Table 3](#).

Form Factor

- PCIe standup: PCI Express short, low-profile card: 167.65mm x 68.90mm (6.60in. x 2.71in.)
- OCP 3.0: Complies with Open Compute Project (OCP) 3.0

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 [Table 2](#) and [Table 3](#).

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/7 |
| 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
- EN IEC 62368 2nd, 3rd Edition
- 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
- EN55035
- 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



QL41132HLRJ




QL41132HQRJ


Table 2. QL41000HLRJ/HQRJ Series Adapters Features

| Adapter Name | QL41132HLRJ | QL41132HQRJ |
|--|---------------------|---------------------|
| General Specs | | |
| Ports | 2 | 2 |
| Port Speeds | 1, 10 | 1, 10 |
| Connectors | BASE-T ¹ | BASE-T ¹ |
| Form Factor | PCIe standup | OCP 3.0 |
| Media | RJ-45 CAT | RJ-45 CAT |
| IEEE 802.3az (EEE) | ✓ | ✓ |
| SmartAN Mode | — | — |
| Storage | | |
| Universal RDMA (RoCE/RoCEv2/iWARP) | ✓ | ✓ |
| NVME-oF over TCP/RDMA | ✓ | ✓ |
| Virtualization and Cloud | | |
| Concurrent SR-IOV/NPAR | ✓ | ✓ |
| DPDK | ✓ | ✓ |
| Flow Filtering | ✓ | ✓ |
| Tunneling Offload (VXLAN/GENEVE/NVGRE/GRE) | ✓ | ✓ |
| Physical Specifications | | |
| Cooling Requirements (LFM/°C) | 150/55 | 225/55 |

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



QL41132HLCU
QL41232HLCU



QL41134HLCU



QL41132HVCU
QL41232HVCU

Table 3. QL41000HLCU/HVCU Series Adapters Features

| Adapter Name | QL41132HLCU | QL41134HLCU | QL41132HVCU | QL41232HLCU | QL41232HVCU |
|--|------------------|------------------|------------------|------------------|------------------|
| General Specs | | | | | |
| Ports | 2 | 4 | 2 | 2 | 2 |
| Port Speeds | 10 | 10 | 10 | 10, 25 | 10, 25 |
| Connectors | SFP+ | SFP+ | SFP+ | SFP+, SFP28 | SFP+, SFP28 |
| Form Factor | PCIe standup | PCIe standup | OCP 3.0 | PCIe standup | OCP 3.0 |
| Media | 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 | ✓ | ✓ | ✓ | ✓ | ✓ |
| Virtualization and Cloud | | | | | |
| Concurrent SR-IOV/NPAR | ✓ | ✓ | ✓ | ✓ | ✓ |
| DPDK | ✓ | ✓ | ✓ | ✓ | ✓ |
| Flow Filtering | ✓ | ✓ | ✓ | ✓ | ✓ |
| Tunneling Offload (VXLAN/GENEVE/NVGRE/GRE) | ✓ | ✓ | ✓ | ✓ | ✓ |
| Physical Specifications | | | | | |
| Cooling Requirements (LFM/°C) ¹ | 100/55 | 150/55 | 100/55 | 100/55 | 100/55 |

1. Cooling requirements are with DAC cables and 85C rated optics; or assume a thermal path from the SFP to the chassis.

Table 4. Ordering Information

| Model Name | HPE Part Number | Description | Form Factor | Protocols |
|-------------|-----------------|--|-------------|------------------------|
| QL41132HLCU | P21933-B21 | Marvell QL41132HLCU Ethernet 10Gb 2-port SFP+ Adapter for HPE | MD2 | L2, RoCE/RocEv2, iWARP |
| QL41134HLCU | P10094-B21 | Marvell QL41134HLCU Ethernet 10Gb 4-port SFP+ Adapter for HPE | MD2 | L2, RoCE/RocEv2, iWARP |
| QL41132HQCU | P08452-B21 | Marvell QL41132HQCU Ethernet 10Gb 2-port SFP+ OCP3 Adapter for HPE | OCP 3.0 | L2, RoCE/RocEv2, iWARP |
| QL41132HLRJ | P08437-B21 | Marvell QL41132HLRJ Ethernet 10Gb 2-port BASE-T Adapter for HPE | MD2 | L2, RoCE/RocEv2, iWARP |
| QL41132HQRJ | P10103-B21 | Marvell QL41132HQRJ Ethernet 10Gb 2-port BASE-T OCP3 Adapter for HPE | OCP 3.0 | L2, RoCE/RocEv2, iWARP |
| QL41232HLCU | P22702-B21 | Marvell QL41232HLCU Ethernet 10/25Gb 2-port SFP28 Adapter for HPE | MD2 | L2, RoCE/RocEv2, iWARP |
| QL41232HQCU | P10118-B21 | Marvell QL41232HQCU Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE | OCP 3.0 | L2, RoCE/RocEv2, iWARP |

All adapters support adaptive voltage scaling (AVS).

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/xalflardzafh32cfvi0z/> for a list of cables and optics that have been tested by Marvell and its partners.

10BGASE-T variants ship with RJ-45 connectors. Intended for use with twisted pair copper cabling (not included).



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 for a complete list of Marvell trademarks. Other names and brands may be claimed as the property of others.