


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.

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.

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)
- 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 QConvergeConsole® (QCC) GUI, 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
- 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: 1/10/25G (1G on 10GBASE-T adapters only)
- 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)

- Data center bridging (DCB) *continued*
 - 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. 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.
- 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 can converge storage and networking I/O by deploying OS-based software iSCSI initiators over their 10GBASE-T and optical or DAC connections.

HIGH-DENSITY SERVER VIRTUALIZATION

The latest hypervisors and multicore systems use several technologies to increase the scale of virtualization. The 41000 Series Adapters adapters support:

- VMware NetQueue
- Windows Hyper-V® Virtual Machine Queue (VMQ)
- Linux Multiqueue
- Windows Hyper-V, Linux Kernel-based Virtual Machine (KVM), and VMware ESXi™ SR-IOV

These features provide ultimate flexibility, quality of service (QoS), and optimized host and virtual machine (VM) performance while providing full bandwidth per port.

Public and private cloud virtualized server farms can now achieve 2.5 times the VM density for the best price and VM ratio.

WIRE-SPEED NETWORK VIRTUALIZATION

Enterprise-class data centers can be scaled using overlay networks to carry VM traffic over a logical tunnel using NVGRE, GRE, VXLAN, and GENEVE. Although overlay networks can resolve virtual LAN (vLAN) limitations, native stateless offloading engines are bypassed, which places a higher load on the system's CPU. The 41000 Series Adapters efficiently handle this load with advanced NVGRE, GRE, VXLAN, and GENEVE stateless offload engines that access the overlay protocol headers. This access enables traditional stateless offloads of encapsulated traffic with native-level performance in the network. Additionally, the 41000 Series Adapters support VMware NSX® and Open vSwitch (OVS™).

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 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)—Azure Stack Hyper Converged Infrastructure (HCI), Windows Live Migration, 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

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)
- 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, v0.90

Ethernet

Throughput

- 10/25Gbps line rate across all ports
- Dual and quad port 10Gbps and dual port 25Gbps
- 1G/10G/25G Auto Negotiation
 - 1G on 10GBASE-T adapters only

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)/VMMQ/vRSS/RSSv2, Linux Multiqueue, and Virtual Machine Device queues (VMDq)
- DDPK
- 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*) (*25G Ethernet*)
 - (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 802.3by Clause 91 (RS-FEC)

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

Operating System Support

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

Packaging

Ports

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

Form Factor

- PCIe standup: 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

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

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%

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

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



QL41132HLRJ



QL41132HQRJ



QL41132HLCU
QL41232HLCU



QL41134HLCU



QL41132HQCU
QL41232HQCU

Table 2. Features

Adapter Name	QL41132HLRJ	QL41132HQRJ	QL41132HLCU	QL41134HLCU	QL41132HQCU	QL41232HLCU	QL41232HQCU
General Specs							
Ports	2	2	2	4	2	2	2
Port Speeds	1, 10	1, 10	1, 10	1, 10	1, 10	1, 10, 25	1, 10, 25
Connectors	BASE-T ¹	BASE-T ¹	SFP+	SFP+	SFP+	SFP+, SFP28	SFP+, SFP28
Form Factor	PCIe standup	OCP 3.0	PCIe standup	PCIe standup	OCP 3.0	PCIe standup	OCP 3.0
Media	RJ-45 CAT	RJ-45 CAT	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							
DPDK	✓	✓	✓	✓	✓	✓	✓
Flow Filtering	✓	✓	✓	✓	✓	✓	✓
Tunneling Offload (VXLAN/GENEVE/NVGRE/GRE)	✓	✓	✓	✓	✓	✓	✓
Physical Specifications							
Cooling Requirements (LFM/°C)	150/55	280/45	100/55	150/55	200/45	100/55	220/45

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

Table 3. Ordering Information

Model Name	HPE Part Number	Description	Form Factor	Platform
QL41132HLCU	P21933-B21	HPE Ethernet 10Gb 2-port SFP+ QL41132HLCU Adapter	MD2	L2, RoCE/RocEv2, iWARP
QL41134HLCU	P10094-B21	HPE Ethernet 10GbE 4-port SFP+ QL41134HLCU Adapter	MD2	L2, RoCE/RocEv2, iWARP
QL41132HQCUC	P08452-B21	HPE Ethernet 10Gb 2-port SFP+ QL41132HQCUC OCP3 Adapter	OCP 3.0	L2, RoCE/RocEv2, iWARP
QL41132HLRJ	P08437-B21	HPE Ethernet 10Gb 2-port BASE-T QL41132HLRJ Adapter	MD2	L2, RoCE/RocEv2, iWARP
QL41132HQRJ	P10103-B21	HPE Ethernet 10Gb 2-port BASE-T QL41132HQRJ OCP3 Adapter	OCP 3.0	L2, RoCE/RocEv2, iWARP
QL41232HLCU	P22702-B21	HPE Ethernet 10/25Gb 2-port SFP28 QL41232HLCU Adapter	MD2	L2, RoCE/RocEv2, iWARP
QL41232HQCUC	P10118-B21	HPE Ethernet 10/25Gb 2-port SFP28 QL41232HQCUC OCP3 Adapter	OCP 3.0	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, QL41132HLCU-BK.

To order a single pack, add *-SP* to the end of the part number; for example, QL41132HLCU-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/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).



Marvell_41000 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.

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.