Marvell® FastLinQ® 41000 Series

Dual-Port, OCP 3.0 Ethernet Network Adapters for Dell®

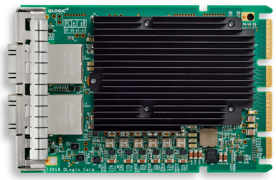

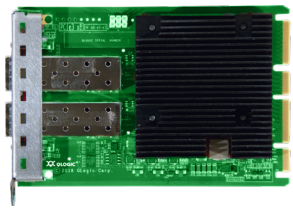

- Dual 10GbE and 25GbE Adapter options
- Delivers full line-rate 10/25GbE performance across dual ports
- Universal RDMA—Delivers the choice and flexibility with concurrent support for RoCE, RoCEv2, and iWARP technologies
- Secure firmware update process with private/public key encryption technology protects against rogue firmware.
- 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 FastLinQ 41000 Series Adapters with Universal Remote Direct Memory Access (RDMA)—available in 10GBASE-T (RJ45), 10-Gigabit Ethernet (GbE) SFP+, and 25GbE SFP28 form factors for Dell PowerEdge® Open Compute Project (OCP) 3.0 Servers—support LAN (TCP/IP) traffic at 10/25GbE line-rate speeds. The 41000 Series provides extremely low host CPU usage by enabling full stateless offloads to meet the performance requirements of the most demanding enterprise applications.

The FastLinQ 41000 Series leverage Marvell's 15+ years of expertise in Ethernet, providing the highest levels of performance, efficiency, and scalability for enterprise data centers.

For more effective use of the 10/25GbE bandwidth, the 41000 Series 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.

- 10GBASE-T version provides low-cost and easy-to-install RJ45 connectivity that is compatible with existing 1GbE
- FastLinQ SmartAN™ for simplified connectivity with switches without user intervention

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 (SW L2 only)
  - 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 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)
- 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

Accelerate Any Network With Universal RDMA Offload

The FastLinQ 41000 Series Adapters support RoCE and iWARP RDMA protocols 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, 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.

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, direct-attach copper (DAC) cables, and active optical cables (AOC), 41000 Series Adapters are the ideal choice for migrating multiple 1GbE network connections to consolidated 10/25GbE.
- Marvell 41000 Series deliver converge storage and networking I/O by deploying an OS-based software iSCSI initiator solution that delivers maximum performance over their 10BASE-T and 10G/25G optical or DAC connections.

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 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) provides vCenter GUI, ESXCLI Plug-ins, and OpenStack integration. QLogic Control Suite (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 Telco Network Function Virtualization (NFV) Workloads

In addition to OpenStack, the Marvell FastLinQ 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. Marvell technology is integrated into the Data Plane Development Kit (DPDK) and can deliver up to 38 million packets per second to host the most demanding NFV workloads.

Also, the FastLinQ 41000 Series Adapters support the NSX-T/N-VDS Enhanced data path/Network Stack (ENS) polling mode driver (QeDeNTV_ens) for NFV workloads on VMware ESXi 6.7.

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 rogue firmware from altering the code running on the adapters.

Host Bus Interface

Bus Interface

- PCI Express® (PCIe) Gen 3 x8

Host Interrupts

- MSI-X supports independent queues

I/O Virtualization

- SR-IOV: up to 192 virtual functions
- Switch-independent NPar/NParEP, up to 16 physical functions
- GRE and NVGRE packet encapsulation offloads
- VXLAN packet encapsulation offloads
- GENEVE packet encapsulation 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 rev. 1.1

Ethernet

Throughput

- 10Gbps line rate per port in 10GbE mode (QL41132HQCU-DE/QL41132HQRJ-DE) and 25Gbps line rate per port in 25GbE mode (QL41232HQCU-DE)
- 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, RSSv2, and TSS
- DPDK
- Enhanced Network Stack for ESXi
- 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)

Ethernet (continued)

Non-offloaded Storage over Ethernet

- iSCSI using OS-based software initiators

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) 802.3-201 8 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*
- RFQs
  - *IPv4 (RFQ 791)*
  - *IPv6 (RFQ 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 (continued)

RDMA Use Cases

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

Tools and Utilities

Management Tools and Device Utilities

- QCS Command Line Interface (CLI) for Linux and Windows
- Plug-in for vSphere (GUI), and ESXCLI plug-in for VMware
- QCC PowerKit (Windows PowerShell®) cmdlets and RESTful APIs for Linux, VMware, and Windows
- Preboot UEFI HII system BIOS device configuration pages
- Native OS management tools for networking
- SNMP support for the Integrated Dell Remote Access Controller (iDRAC) and other application tools

Boot Support

- Pre-execution environment (PXE) 2.0
- UEFI
- iSCSI remote boot (SW L2 only)

Operating System Support

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

Forward Error Correction (FEC)

- FireCode or BASE-R IEEE 802.3-2018 Clause 74 or FC-FEC
- Reed Solomon IEEE 802.3-2018 Clause 91 or RS-FEC

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

Picture may not be representative of the final shipping product.

Physical Specifications

Ports

- Dual 10Gbps Ethernet SFP+ cages (QL41132HQCU-DE)
- Dual 10GBASE-T (QL41132HQRJ-DE)
- Dual 25Gb SFP+ cages (QL41232HQCU-DE)

Form Factor

- OCP 3.0 Small Form Factor (SFF) rev 1.1

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 on page 8.

Humidity (Relative, Non-condensing)

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

Compliance

- RoHS complaint

Connectivity

Table 1. Cable Distance—QL41132HQCU-DE and QL41232HQCU-DE

Rate	Cable and Maximum Distance (m)		
	DAC	SR FOC	AOC
10G[1, 2]	7	400 OM4/5 300 OM3	20
25G[2]	5	100 OM4 70 OM3	20

1. QL41132HQCU-DE, QL41232HQCU-DE
2. QL41232HQCU-DE

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

Table 2. Cable Distance—QL41132HQRJ-DE

Rate	Cable and Maximum Distance (m)	
	Cat 6	Cat 6a/7
1GBASE-T	100	100
10GBASE-T	37 to 55	100

Approvals—Safety

US and Canada

- UL 60950-1
- CSA C22.2
- EN62368

Europe

- TUV EN60950-1
- TUV IEC 60950-1
- TUV EN62368
- IEC 62368 2nd 3rd -1 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

United Kingdom

- UKCA

Japan

- VCCI: Class A

New Zealand and Australia

- AS/NZS: Class A

Korea

- KC-RRA Class A

Taiwan

- BSMI CNS 13438

Table 3. Features

Adapter Name	QL41132HQCU-DE	QL41132HQRJ-DE	QL41232HQCU-DE
General Specs			
Ports	2	2	2
Port Speeds	10	1/10	10/25
Connectors	SFP+	BASE-T	SFP+, SFP28
Form Factor	OCP 3.0	OCP 3.0	OCP 3.0
Media	DAC, Optics, AOC	RJ-45	DAC, Optics, AOC
802.3az (EEE)	—	√	—
Advanced Configuration and Power Interface (ACPI), v2.0	√	√	√
SmartAN™ Mode	—	—	√
Storage			
Universal RDMA (RoCE/RoCEv2/iWARP)	√	√	√
Virtualization and Cloud			
Concurrent SR-IOV/NPAR	√	√	√
DPDK	√	√	√
Flow Filtering	√	√	√
Tunneling Offload (VXLAN/ GENEVE/NVGRE/GRE)	√	√	√
Physical Specifications			
Operating Temperature	0°C to 55°C 32°F to 131°F	0°C to 55°C 32°F to 131°F	0°C to 55°C 32°F to 131°F
Cooling Requirements (LFM/°C)	100 LFM at 55°C with DAC cable or 85C optics	225LFM at 55°C	100LFM at 55°C with DAC cable or 85C SR optics

Table 4. Ordering Information

Part Number	Factory Install SKU	Customer Install SKU	Description	Form Factor
QL41132HQCU-DE	540-BCNX	540-BCOX	2x 10GbE SFP+	OCP NIC 3.0
QL41232HQCU-DE	540-BCNU	540-BCOY	2x10/25GbE SFP28	OCP NIC 3.0
QL41132HQRJ-DE	540-BCNY	540-BCOW	2x 10GbE Base-T	OCP NIC 3.0

All adapters support adaptive voltage scaling (AVS).

Twisted pair cabling, DAC cables, SR/LR optics are not included. See https://www.marvell.com/content/dam/marvell/en/public-collateral/ethernet-adaptersandcontrollers/marvell-ethernet-adapters-fastlinq-41000-interoperability-matrix-2019-06.pdf for a list of cables and optics that have been tested by Marvell and its partners. The Marvell SFP+ optic is not supported by QL41132HQRJ-DE.

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