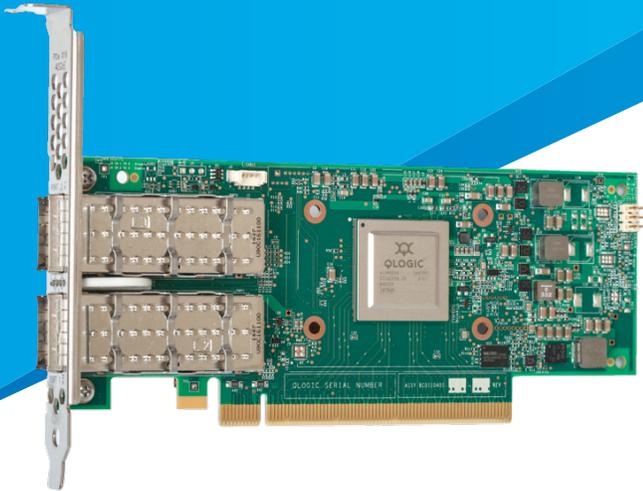# FastLinQ QL45412HLCU-CI

40GbE Intelligent Ethernet Adapter

Note: Picture may not be representative of the final shipping product

- Fully featured 40GbE adapter delivers a better price-per-performance ratio versus 10GbE

- Increase VM density and accelerate multitenant networks with full offload for tunneling protocols

- Build the most powerful scale-out storage systems with Cavium's unique support of universal RDMA technologies (RoCE, RoCEv2, and iWARP)

- Accelerate the most demanding telco NFV workloads with the Cavium DPDK high-speed packet processing engine

- Orchestrate and manage hyperscale OpenStack® deployments with Cavium QConvergeConsole® cloud-enabled management framework

OVERVIEW

The Cavium QL45412HLCU-CI dual-port Intelligent Ethernet Adapter leverages Cavium's seventh-generation technology to deliver true 40Gbps Ethernet performance. Optimized for use across enterprises, managed service providers, and large public and scalable private cloud deployments, the QL45412HLCU-CI enables organizations to achieve new levels of performance in physical, virtual, and cloud environments.

The QL45412HLCU-CI 40GbE Adapter delivers advanced features, including:

- Cutting-edge server virtualization technologies—single-root I/O virtualization (SR-IOV) and NIC partitioning (NPAR)

- Network virtualization—offloads for VXLAN, GENEVE, and NVGRE

- Multiple, concurrent RDMA technologies—RDMA over Converged Ethernet (RoCE), RoCEv2, iSCSI Extensions for RDMA (iSER), and is extensible to support iWARP

REDUCE CAPITAL EXPENDITURES (CAPEX) AND OPERATIONAL EXPENDITURES (OPEX)

Cavium QL45412HLCU-CI 40GbE technology delivers better price-per-gigabit ratio versus 10GbE. This technology enables cloud providers and large-scale data center operators to reduce operating expenses while continuing to scale their network of server and storage nodes to meet

increasing demands of the future.

ACCELERATE ANY NETWORK WITH UNIVERSAL RDMA OFFLOAD

Cavium QL45412HLCU-CI 40GbE technology supports RoCE acceleration to deliver low latency, low CPU utilization, and high performance on Windows Server® Message Block (SMB) Direct 3.0 and 3.02. The QL45412HLCU-CI 40GbE Adapter has the unique capability to deliver universal RDMA that provides customers with the choice of a low-latency interconnect that can best utilize the potential of emerging storage technologies such as Non-Volatile Memory Express over Fabric (NVMe-oF) and Network File System over RDMA (NFSoRDMA). Cavium's cutting-edge offloading technology increases cluster efficiency and improves scalability.

HIGH-DENSITY SERVER VIRTUALIZATION

The latest hypervisors and multicore systems use several technologies to increase the scale of virtualization. The QL45412HLCU-CI 40GbE adapter supports:

- VMware® NetQueue

- Windows® Hyper-V® Virtual Machine Queue (VMQ)

- Linux® Multiqueue

- Windows, Linux, and VMware switch-independent NPAR

- Windows Hyper-V, Linux Kernel-based Virtual Machine, and VMware ESXi™ SR-IOV

These features provide ultimate flexibility, QoS, and optimized host and VM performance while providing full 40Gbps bandwidth per port. Public and private cloud virtualized server farms can now achieve four 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, VXLAN, and GENEVE. Although overlay networks can resolve VLAN limitations, native stateless offloading engines are bypassed, which places a higher load on the system's CPU. Cavium QL45412HLCU-CI 40GbE technology efficiently handles this load with advanced NVGRE, VXLAN, and GENEVE stateless offloading engines that access the overlay protocol headers. This access enables traditional stateless offloads of encapsulated traffic with native-level performance in the network. Additionally, Cavium QL45412HLCU-CI 40GbE technology supports VMware NSX® and Open vSwitch.

HYPERSCALE ORCHESTRATION WITH OPENSTACK
Cavium QL45412HLCU-CI 40GbE technology supports the OpenStack open source infrastructure for deploying and orchestrating public, private, and hybrid cloud computing platforms. It provides for both networking and storage services (block, file, and object) for iSER. These platforms allow providers to rapidly and horizontally scale VMs over their entire, diverse,

and widely spread network architecture to meet the real-time needs of their customers. The integrated, multiprotocol management utility, QConvergeConsole, provides breakthrough features that allow customers to visualize the OpenStack-orchestrated data center using autodiscovery technology.

ACCELERATE TELCO NETWORK FUNCTION VIRTUALIZATION (NFV) WORKLOADS
In addition to OpenStack, Cavium QL45412HLCU-CI 40GbE technology supports NFV that 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, which reduces CAPEX and OPEX, and enhances business and network services agility. Cavium 40GbE technology is integrated into the Data Plane Development Kit (DPDK) to host the most demanding NFV workloads.

TRUSTED, RELIABLE, AND INTEROPERABLE
Cavium QL45412HLCU-CI 40GbE technology adheres to standards that ensure interoperability with a wide range of network solutions. Using advanced Cavium technologies based on mature software stacks, enterprise-class data centers can confidently deploy reliable, high-performance networks.

Host Bus Interface Specifications

Bus Interface
- PCIe® Gen3x16, PCIe Gen2x16, PCIe Gen1x16

Host Interrupts
- INTx and MSI-X

I/O Virtualization
- SR-IOV (up to 240 virtual functions)
- NPAR (up to 16 physical functions)

Compliance
- *PCI Express Base Specification,* rev. 3.1
- *PCI Express Card Electromechanical Specification,* rev. 3.0
- *PCI Bus Power Management Interface Specification,* rev. 1.2

Ethernet Specifications

Throughput
- 40Gbps line rate per port

Ethernet Frame
- Standard MTU sizes and jumbo frames up to 9,600 bytes

Stateless Offload
- IP, TCP, and user datagram protocol (UDP) checksum offloads
- TCP segmentation offload (TSO)
- Large send offload (LSO)
- Giant send offload (GSO)
- Large receive offload (LRO)
  - LRO (Linux)
  - Receive segment coalescing (RSC) (Windows)
- Receive side scaling (RSS)
- Transmit side scaling (TSS)
- Interrupt coalescing
- VMware NetQueue, Microsoft® Hyper-V VMQ (up to 512 queues), and Linux Multiqueue

Ethernet Specifications (*continued*)

Network Virtualization
- GRE
- NVGRE
- VXLAN
- GENEVE

Compliance
- IEEE Specifications
  - *40GBASE-CR4 (Direct Attach Copper)*
  - *40GBASE-SR4 (Multimode Fiber)*
  - *802.3-2012 (40Gb Ethernet and 10Gb Ethernet and Ethernet Flow Control)*
  - *802.1q (VLAN)*
  - *802.1AX- (Link Aggregation)*
  - *802.1p (Priority Encoding)*
  - *IPv4 (RFQ 791)*
  - *IPv6 (RFC 2460)*
  - *802.1Qbb (Priority-Based Flow Control)*
  - *802.1Qaz (DCBX/Enhanced Transmission Selection)*
  - *802.1Qau (Congestion Notification)*
  - *1588-2002 PTPv1 (Precision Time Protocol)*
  - *1588-2008 PTPv2*

RDMA Specifications

Universal RDMA
- RoCE
- RoCEv2
- Storage over RDMA: iSER, SMB Direct, and NVMe™ over Fabrics
- NFSoRDMA

Forward Error Correction (FEC) Support

- FC-FEC

Tools and Utilities

Management Tools and Device Utilities
- QLogic Control Suite integrated network adapter management utility (CLI) for Linux and Windows
- QCC PowerKit PowerShell cmdlets for Linux and Windows
- QConvergeConsole integrated network management utility (GUI) for Linux and Windows
- QConvergeConsole Plug-ins for vSphere® (GUI) and ESXCLI plug-in for VMware
- Native OS management tools for networking

Boot Support
- Unified extensible firmware interface (UEFI)
- Pre-execution environment (PXE)
- Software iSCSI boot (iBFT)

Microsoft SDDC-Premium Certified for Azure Stack Hybrid Cloud and Storage Spaces Direct Solutions
- Routable RoCEv2 for SMB direct
- RDMA over virtual switch traffic
- Data center bridging (DCB) for QoS
  - Enhanced transmission selection (ETS)
  - Priority flow control (PFC)
- Hyper-V network virtualization (HNV) encapsulation Offloads for VXLAN and NVGRE
- Virtual machine multiqueues (VMMQ)
- SR-IOV
- Hyper-V Switch Embedded Teaming (SET)

Operating Systems
- For the latest applicable OS information, see http://driverdownloads.qlogic.com

Physical Specifications

Ports
- Dual 40Gbps Ethernet QSFP+ optical or DAC ports

Form Factor
- Low profile PCIe card (6.6in. × 2.54in.)

Environment and Equipment Specifications

Temperature
- Operating: 0°C to 55°C (32°F to 131°F)
- Storage: –20°C to 70°C (–4°F to 158°F)

Humidity
- Operating: 10% to 90%
- Storage: 5% to 95%

Maximum Cable Distances
- 100m OM3 MMF or 150m OM4 MFF
- 300m OM3 (if enhanced QSFP optical modules are installed on both ends of the link)
- 7m DAC

Agency Approvals—Safety

US and Canada
- UL 60950-1
- CSA C22.2

Europe
- TUV EN60950-1
- TUV IEC 60950-1
- CB Certified

Agency Approvals[1]—EMI and EMC (Class A)

US and Canada
- FCC Rules, CFR Title 47, Part 15, Subpart Class A
- Industry Canada, ICES-003: Class A

Europe
- EN55022
- EN55024
- EN61000-3-2
- EN61000-3-3

Japan
- VCCI: Class A

New Zealand and Australia
- AS/NZS: Class A

Korea
- KC-RRA Class A

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

1   Agency approvals are preliminary at the time of initial release.

Taiwan
• BSMI CNS 13438

Compliance
• RoHS compliant

Ordering Information

QL45412HLCU-CI-BK Dual-port 40Gb Ethernet
Network Adapter Card
• Cisco product ID: UCSC-PCIE-QD40GF and UCSC-
  PCIE-QD40GF=
• Ships in a bulk-packed box with a standard-height
  bracket installed
• Ships without SR optical transceivers installed

Follow us: