

Oracle Quad 10 Gb Ethernet Adapter or Oracle Dual 40 Gb Ethernet Adapter

Introduction

Oracle's latest Ethernet Adapter the Oracle Quad 10 Gb Ethernet Adapter or Oracle Dual 40 Gb Ethernet Adapter is the latest high-bandwidth, low-power network interface card (NIC) based on the Intel® Ethernet Controller XL710 controller. This dual-port QSFP+ 3.0 x8 low-profile adapter is designed to provide multi-port network connectivity in data centers. Today's cloud-enabled software-defined data centers require higher networking speeds and higher performance for mission-critical applications, and Oracle Quad 10 Gb Ethernet Adapter or Oracle Dual 40 Gb Ethernet Adapter addresses these needs with two QSFP+ interfaces providing up to 80 Gb/sec bandwidth.



Figure 1: Oracle Quad 10 Gb Ethernet Adapter or Oracle Dual 40 Gb Ethernet Adapter

Frequently Asked Questions

Q: What is new in Oracle Quad 10 Gb Ethernet Adapter or Oracle Dual 40 Gb Ethernet Adapter?

A: * Advanced single root I/O virtualization SR-IOV capabilities with up to 16 physical functions (PFs) and up to 128 virtual functions (VFs)

* Support for network overlay using Network Virtualization using Generic Routing Encapsulation (NVGRE) or Virtual Extensible LAN (VXLAN)

* Sophisticated packet header parsing with up to 64 different packet types

* Simplified network convergence by consolidating LAN and SAN (iSCSI)

Q: What are the marketing part numbers?

A: **7114148** Oracle Quad 10 Gb Ethernet Adapter or Oracle Dual 40 Gb Ethernet Adapter (for factory installation)

7114134 Oracle Quad 10 Gb Ethernet Adapter or Oracle Dual 40 Gb Ethernet Adapter (X-option)

Q: What is the adapter default port mode?

A: Oracle Quad 10 Gb Ethernet Adapter or Oracle Dual 40 Gb Ethernet Adapter comes with a factory default port mode of 4x 10 GbE with only the top QSFP port enabled. However, with the software utility download from MyOracleSupport, the ports can be dynamically configured to 2x 40 GbE or 4x 10 GbE. After the software utility is invoked, PCIe reset is required to transition to the new mode completely.

Q: What are the cabling options for the modes?

A: Currently only copper cabling can be used to connect this adapter. Here is a detailed explanation of the required cables:

Copper solution (for 2x 40 Gb mode):

QSFP to QSFP copper cables (available in 1/2/3/5 meters). One end of the QSFP cable plugs into a QSFP port on the NIC, and the other end goes into a QSFP port on a switch or other device.

Copper solution (for 4x10 Gb mode):

QSFP copper splitter cables (available in 1/3/5 meters). The QSFP end of the cable plugs into the QSFP port on the NIC, and the other four pigtailed ends can plug into four

different 10 Gb SFP+ ports on a switch or any other device.

Copper solution does not require QSFP transceivers.

Q: Which transceiver is supported with this adapter?

A: Currently there is no transceiver that is supported. In the near future the supported transceiver will be the QSFP optical transceiver option:

(X) 2124A-N QSFP optical SR (supports multimode cables with MPO connector up to 100 m)

Q: Does this adapter support NetQueue?

A: Yes. This adapter supports VMDq for VMware NetQueue and Microsoft's DVMQ.

Q: What network overlay technologies are supported?

A: This adapter supports NVGRE, VXLAN, GENEVE, IPinGRE, and MACi n UDP overlay technologies.

Q: Is FCoE supported on this adapter?

A: No. This adapter does not support FCoE.

Q: Are jumbo frames supported?

A: Jumbo frames (9,706 bytes) are supported.

Q: Does this card support 802.1 P quality of service (QoS) tagging?

A: Yes. 802.1p QoS tagging is supported. Supports eight VLAN filters per port. Filtering is ignoring the *CFI* bit and the 802.1P priority bits.

Q: Does it support TCP offload?

A: Yes. TCP segmentation offload (TSO)/large receive offload (LRO) are supported for improved performance.

Q: Does the adapter support PXE boot?

A: Yes. PXE and iSCSI boot capabilities are supported.

Q: Which hashing algorithms are supported?

A: Toeplitz and symmetric receive side scaling (RSS), and XOR.

Q: What are the queue selection features for forwarding traffic from VSIs to receive engines?

A: It could be perfect match based on L2/L3 fields or hash via RSS or default queue.

Q: What are advanced traffic steering capabilities?

A: This adapter supports Intel Ethernet Flow Director with 8k perfect match filters stored on the die. The adapter also supports up to 64 different packet types (L2/IPV4/IPV6/TCP/UDP/SCTP) with filtering on flexible offsets.

Q: What are some of the performance optimizations that are supported?

A: Packet checksum offloads (UDP, TCP, IP, and Stream Control Transmission Protocol)

TCP/UDP segmentation offload, also referred to as large send offload (LSO) for transmitted IPv4 and IPv6 packets. The maximum size of a TSO operation is 256 KB.

Received packet header splitting

RSS: PF instances of the RSS logic implement 256 entry indirection tables, supporting up to 64 RQs/CPU. VF instances of the RSS logic implement 64 entry indirection tables, supporting up to 16 RQs/CPU.

Cabling Options

QSFP to QSFP Direct-Attach Passive Copper Cable Options

(QSFP transceivers not required)

A: **X2121A-1M-N** QSFP to QSFP passive copper cable, 1 meter

X2121A-3M-N QSFP to QSFP passive copper cable, 3 meters

X2121A-5M-N QSFP to QSFP passive copper cable, 5 meters

QSFP to 4 SFP+ Passive Copper Splitter Cable Options

X2125A-1M-N Copper splitter cable assembly: 1 meter, QSFP+ to 4 SFP+

X2125A-3M-N Copper splitter cable assembly: 3 meters, QSFP+ to 4 SFP+

X2125A-5M-N Copper splitter cable assembly: 5 meters, QSFP+ to 4 SFP+

Oracle Corporation, World Headquarters

500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries

Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US

blogs.oracle.com/blogs

facebook.com/oracle

twitter.com/oracle

oracle.com



Integrated Cloud Applications & Platform Services

Copyright © 2020, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0316