

Chelsio 10 GbE Adapters for IBM Power Systems

IBM Power at-a-glance guide

Chelsio 10 GbE adapters for IBM Power Systems are high performance, 10 Gb Ethernet adapters with PCI Express host bus interfaces optimized for virtualization, high performance computing, and storage applications. The third generation technology from Chelsio used in these adapters provides the highest 10 GbE performance available and dramatically lowers host system's CPU communications' impact.

With on-board hardware that off-loads TCP/IP, iSCSI, and iWARP RDMA processing from its host system, single port adapters frees up host CPU cycles for useful applications, which results in increased bandwidth and lower latency. This combination makes it practical to converge other networks that traditionally used niche technologies onto 10 GbE Ethernet adapters. High band-width and extremely low latency make 10 GbE Ethernet adapters with protocol offload the best technology for high-performance cluster computing (HPCC) fabrics.

Figure 1 shows the Chelsio single port 10 GbE adapters for IBM Power Systems.



Figure 1. The 10 Gigabit Ethernet-SR PCI Express Adapter (left) and 10 Gigabit Ethernet-CX4 PCI Express Adapter (right)

Did you know

Chelsio delivers on the promise of the *unified wire*, enabling the convergence of server networking, storage networking, and cluster computing interconnects onto a single platform and a single fabric.

Part number information

Table 1. Ordering part number and feature code

Description	IBM part number	IBM feature code	Chelsio part number
10 Gigabit Ethernet-SR PCI Express Adapter	46K7897	5769	S310E-SR+
10 Gigabit Ethernet-CX4 PCI Express Adapter	46K7899	5732	S310E-CXA

Features

The 10 Gigabit Ethernet-SR and CX4 PCI Express Adapters are low-profile, high performance adapters. The SR adapter uses a LC Duplex type connector that is capable of transferring data a distance of 300 m over MMF-850 nm fiber cable and conforms to the 10GBASE-SR specification. The CX4 adapter uses a CX4 type connector that is capable of transferring data a distance of 15 m over copper cabling and conforms to the 10GBASE-CX4 specification.

The adapters have the following features:

- Third generation technology (T3) from Chelsio
- The Unified Wire interconnect solution for server networking, storage networking, and clustering on a single platform
- Very low latency Ethernet
- Reduced host CPU utilization by up to 90% compared to NICs without full offload capabilities
- PCI Express 1.1 x8 host bus interface
- Line-rate 10 Gbps full-duplex performance
- Integrated traffic manager, QoS, and virtualization capabilities
- TCP sockets, SCSI, RNIC-PI, kDAPL, and OpenFabrics 1.3 software interfaces
- Ideal fit for servers and NAS systems
- Consolidated LAN, SAN, and cluster networks with TOE, iSCSI, and iWARP
- Enabled high performance NAS systems and Ethernet-based IP SANs
- Deployed Ethernet-only IP SANs throughout the enterprise beyond the data center
- Developed shared-storage systems providing both file- and block-level services
- Deployed Ethernet-only networking for cluster fabric, LAN, and SAN
- Full TCP offload, iSCSI, and iWARP RDMA plus direct data placement (DDP)
- Increased cluster fabric bandwidth and deployed Ethernet-only networking for cluster fabric, LAN, and SAN

Table 2 compares the adapters.

Table 2. Comparison of the adapter cards

Feature	10 Gigabit Ethernet-SR PCI Express Adapter	10 Gigabit Ethernet-CX4 PCI Express Adapter
Feature code	5769	5732
Form factor	Low profile PCIe card	Low profile PCIe card
Ports	1	1
Connector	SFP+ transceiver	CX4
Power consumption	17 W	16 W

Specifications

The Chelsio 10 GbE adapters for IBM Power Systems have the following specifications:

- Network interfaces
 - 10 Gigabit Ethernet-SR PCI Express Adapter: 10GBASE-SR short-reach optics (850 nm)
 - 10 Gigabit Ethernet-CX4 PCI Express Adapter: 10GBASE-CX short-reach copper

The 10 Gigabit Ethernet-SR PCI Express Adapter also supports 10GBASE-LR long-reach optics (1310 nm) with the use of a LR transceiver (optional, available directly from Chelsio using part number SM10G-LR)

- Ethernet
 - IEEE 802.3ae (10 GbE)
 - IEEE 802.1p Priority and 802.1Q VLAN tagging
 - IEEE 802.3x flow control
 - IEEE 802.3ad link aggregation
 - Ether II and 802.3 encapsulated frames
 - Multiple MAC addresses per interface
 - Jumbo Frames up to 9.6 KB
- Stateless off-loads
 - TCP checksum offload for IPv4 and IPv6
 - TCP Segmentation Offload (TSO) for IPv4 and IPv6
 - Large Receive Offload (LRO) for IPv4 and IPv6
 - Large Send Offload (LSO) for IPv4 and IPv6
 - UDP checksum offload for IPv4 and IPv6
 - Receive Side Scaling and packet steering
 - Line rate packet filtering and attack protection

- TCP/IP full offload
 - Full TCP implementation including exceptions
 - High performance even during packet loss
 - Extensive RFC compliance, fully featured stack
 - Zero copy on Tx and Rx
 - Direct Data Placement (DDP)
 - Up to 64 K simultaneous connections capacity
- Integrated traffic manager
 - Multiple Tx and Rx queues with QoS
 - Simultaneous low latency and high bandwidth
 - Per-connection and per-class rate control
 - Packet loss avoidance
- Virtualization and firewall
 - Rule-based packet steering and filtering capability
- iSCSI acceleration
 - Full iSCSI initiator and target mode stack
 - Header and Data Digest (CRC) generation and checking
 - PDU recovery
 - Direct Data Placement (DDP)
- High performance RDMA
 - Ultra-low latency and line rate bandwidth
 - IETF RDDL and RDMAC iWARP compliance
 - APIs: RNIC-PI, kDAPL, and OpenFabrics 1.3

RDMA features are currently only supported by Linux.

Physical specifications

The adapters have the following physical specifications:

- Width: 168 mm (6.60 in.)
- Height: 64 mm (2.52 in.)
- Depth: 18 mm (0.69 in.)

Operating environment

The adapters are supported in the following environment:

- Operating temperature: 0 to 55° C (32° to 131° F)
- Operating humidity: 5 to 95%
- Typical power consumption:
 - 10 Gigabit Ethernet-SR PCI Express Adapter: 17 W
 - 10 Gigabit Ethernet-CX4 PCI Express Adapter: 16 W
- Minimum required airflow: 200 lf/m

Supported servers

The Chelsio 10 GbE adapters for IBM Power Systems are supported in the IBM servers that are listed in Table 3.

Table 3. Supported servers

Adapter	Feature	Power 520 Express	Power 560 Express	Power 570	Power 575	Power 595	Power 750	Power 755	Power 770	Power 780	Smart Cube
10 Gigabit Ethernet-SR PCI Express Adapter	5769	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
10 Gigabit Ethernet-CX4 PCI Express Adapter	5732	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Supported operating systems

The Chelsio adapters support the following operating systems (including support for both TOE and NIC mode, unless otherwise stated):

- Red Hat Enterprise Linux 3 base kernel (NIC only; no TOE support)
- Red Hat Enterprise Linux 4 U3, U4, U5, U6, U7, and U8
- Red Hat Enterprise Linux 5 base, U1, U2, U3, and U4
- SUSE Linux Enterprise Server 10 base, SP1, and SP2
- SUSE Linux Enterprise Server 9 SP3 (NIC only; no TOE support)
- SUSE Linux Enterprise Server 11 base kernel
- Kernel.org linux-2.6.9, 2.6.12, 2.6.16, 2.6.20
- AIX 5.3 (NIC only; no TOE support)
- AIX 6.1 (NIC only; no TOE support)

Note: The supported servers will have other operating system support requirements that must also be met. Check the Software Requirements section of the relevant announcement letter of the server for the specific operating system levels supported.

See the Chelsio driver README file for the latest information about the specific versions and service packs that are supported. This README file is available at the following address:

<http://www.chelsio.com/support>

Related publications

For more information, see these resources:

- Chelsio Adapters driver download, Installation, Configuration, troubleshooting, support, found at the following address:
<http://service.chelsio.com/ibm.html>
- IBM Announcement Letter, found at the following address:
<http://ibm.com/common/ssi/cgi-bin/ssialias?infotype=dd&subtype=ca&&htmlfid=897/ENUS109-305>

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service. IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you. This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk. IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

© Copyright International Business Machines Corporation 2010. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

This document was created or updated on April 14, 2010.

Send us your comments in one of the following ways:

- Use the online **Contact us** review form found at:
ibm.com/redbooks
- Send your comments in an e-mail to:
redbook@us.ibm.com
- Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400 U.S.A.

This document is available online at <http://www.ibm.com/redbooks/abstracts/tips0774.html> .

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®
IBM®
Power Systems™
Redpaper™
Redbooks (logo)®
System x®

The following terms are trademarks of other companies:

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.