




Third-party information brought to you courtesy of IBM.



## Product Brief

**Intel® Ethernet  
SFP+ Dual Port  
Embedded CNA X520-DA2  
for IBM Servers**

### Key Features

- Low cost, low power, 10 Gigabit Ethernet (10GbE) performance for the entire data center
- Twinaxial cabling with Direct Attach SFP+ or SR Fiber Optic connector
- Backward compatibility with existing 1000BASE-T networks simplifies the transition to 10GbE
- Flexible I/O virtualization for port partitioning and quality of service (QoS) of up to 64 virtual ports
- Unified Networking delivering LAN and iSCSI in one low-cost CNA
- Reliable and proven 10 Gigabit Ethernet technology from Intel

# Intel® Ethernet 10G SFP+ DP Embedded CNA X520-DA2

Dual-Port SFP+ Converged Network Adapter for Providing Ultimate Flexibility and Scalability in Virtual and Unified Storage Environments


### Overview

The best-selling Intel X520 dual-port Converged Network Adapters are known for their high performance, low latency, reliability, and flexibility. The addition of the new Intel® Ethernet 10GSFP+ DP Embedded CNA X520-DA2 adapters to the IBM family delivers all of the X520 capabilities in an embedded adapter. The Intel® Ethernet Converged Network Adapter X520 family of adapters delivers a proven, reliable solution for deployments of high density Ethernet for unified 10GbE network connections. Increased I/O bandwidth and support for intelligent offloads allow for scaling performance on platforms using the Intel® Xeon® processor E5/E7 family. These cards deliver 20GbE of bandwidth today for high bandwidth application requirements in an embedded card that does not consume a PCIe slot resource leaving them available for other customer needs.

Powered by the Intel® Ethernet 82599 Ethernet Controller, the Embedded X520-DA2 CNA addresses the demanding needs of the next-generation data center by providing unmatched features for virtualization, flexibility for LAN and SAN networking, and proven, reliable performance. Intel® 82599 controller is the industry standard for 10GbE making it the most popular 10GbE controller on the market.

### Best Choice for Virtualization

Intel leads the industry in virtualization by being the first to provide virtualization for all the major operating systems and working with IBM to implement virtualization not only on the adapter but also on the IBM platform.



### Intel® Virtualization Technology for Connectivity (Intel® VT-c)

Intel® Ethernet Controllers include Intel® Virtualization Technology for connectivity (Intel VT-c) to deliver virtualized I/O performance optimizations and Quality of Service (QoS) features designed directly in to the controller's silicon. Virtual Machine Device Queues (VMDq) and PCI-SIG\* Single Root I/O Virtualization and Sharing (SR-IOV) are two of the core optimization technologies that help reduce I/O bottlenecks and improve the overall server performance. Both VMDq and SR-IOV offload the sorting and queuing functionality to the Ethernet controller which reduces the data processing bottleneck associated with virtualization. Additionally, VMDq and SR-IOV use some of the same built-in QoS feature to balance bandwidth and improved I/O scalability.

### Hypervisor Hardware Assist using Virtual Machine Device Queues (VMDq)

VMDq works in conjunction with VMware\* NetQueue\* or Microsoft\* Virtual Device Queues\* (VMQ) in their respective hypervisors to use the hardware assists on the controller

for traffic steering and Rx/Tx round-robin scheduling for balancing bandwidth allocation across transmit and receive queues. Based on processor and network load, the hypervisor dynamically offloads routing and filtering of network packets to hardware-based receive queues on the controller to reduce CPU cycles.

## Hypervisor BYPASS using SR-IOV

Bypassing the hypervisor and allowing direct hardware access by virtual machines, significantly reduces CPU overhead, reduces latency, and increases network throughput. Most of the current hypervisor releases have been enabled to partition a single physical Ethernet controller into multiple virtual Ethernet controllers that can be used directly by VMs, commonly known as Direct Assignment. The use of these virtual controllers, known as Virtual Functions (VF), enables additional QoS features in the controller's silicon to manage and direct traffic such as traffic isolation, port partitioning with bandwidth allocation and on-chip VF-VF switching.

## Flexible Port Partitioning Support

Intel Flexible Port Partitioning (FPP) is the ability to use SR-IOV Virtual Functions, (which have, up until now, been thought of strictly as a virtualization technology) in a bare-metal (or mixed) Open Source OS. This provides a way to very flexibly and efficiently divide your Ethernet ports.

The PCI-SIG® Single-Root I/O Virtualization and Sharing Technology (SR-IOV) standard defines the partitioning of a single Ethernet port into multiple virtual functions that are seen by an OS as virtual adapters.

These virtual adapters can be used by the Linux host directly and/or assigned to virtual machines. This port partitioning allows administrators to create multiple isolated networks on a single Ethernet port for use in bare-metal and virtualized server deployments.

In a bare-metal Linux server, host processes can be assigned to dedicated network resources to provide traffic isolation. Also, the hardware-based QoS functionality keeps the network connections available for critical traffic, even during heavy traffic contention. These isolation and QoS features can be used even in a 1GbE environment to help ensure that management traffic has access to the host OS.

In a virtualized environment, a direct assignment of a virtual machine (VM) to a virtual adapter reduces the CPU overhead seen when using a software-based network bridge or virtual switch by offloading network traffic management to the Ethernet controller silicon.

Additionally, bandwidth throttling and monitoring capabilities allow for fine tuning of QoS requirements for each of the virtual adapters whether used by the OS or by a VM. All this can be

accomplished with standard 10GbE switches, providing a cost effective solution for users that want to segment and throttle their network traffic.

## Unified Networking and Storage

The family of Intel® Ethernet CNA X520 adapters lowers your data center total cost of ownership (TCO) by providing the ability to route LAN and storage traffic over a single fabric. By enabling IT to unify LAN and storage traffic in a single network, Intel helps reduce costs and complexity, and enhance the efficiency of your virtual infrastructure.

Unified networking is ideal for cloud usage. Discrete data center networks for specific components such as storage and servers are no longer appropriate. As companies move to cloud computing, networks must be simpler to manage while supporting large volumes of traffic. Converged networks, with storage and server traffic running across a unified fabric, do that.

The benefits of such a unified network are simplification, lower TCO due to infrastructure consolidation, and greater flexibility.

In today's dynamic data center environment, you need to simplify as much as possible. Thanks to unified networking from Intel, you can combine all the traffic of multiple data center networks—such as iSCSI, FC, TCP/IP and NAS—using one efficient network fabric. You can also get world-class iSCSI and FCoE performance with native OS initiators and intelligent CPU and networking hardware offloads. Such innovations add up to a simpler, more efficient data center environment.

Intel has lead broad OS and software ecosystem engagements to enable effective protocol processing for storage and LAN traffic through native software storage initiators. Native initiators are:

- **Easy to use**—they have the same management interfaces that are integrated into the Operating system.
- **Reliable**—Validated and backed by the reliability and trust from Operating system vendor. With multiple OS vendors and storage companies validating the storage initiator, the end user benefits with the highest quality solution offered in the market.
- **Cost Effective**—integrated into the Kernel means customers can implement over a standard Ethernet connection and don't have to pay extra for the feature.

## Enabled for Fibre Channel Over Ethernet (FCoE)

FCoE encapsulates Fibre Channel frames over standard Ethernet networks, enabling Fibre Channel to take advantage of 10 GbE networks while preserving its native protocol. The X520 adapters offer FCoE hardware acceleration to provide performance comparable to FC HBAs. The adapters support Data Center Bridging (DCB), also known as Converged Enhanced Ethernet (CEE), which enables customers to configure traffic classes and priorities

to deliver a lossless Ethernet fabric. An Intel Ethernet Converged Network Adapter X520 reduces TCO by eliminating redundant fabrics and saves the cost of expensive FC HBAs and FC switch ports.

## Support for iSCSI

The adapters provide complete support for proven native OS iSCSI initiators as well as iSCSI boot. Historically, CRC32C computation has degraded system performance, but now, with the CRC instruction set included in the latest Intel® Xeon® processors, CRC validation is possible with minimal impact to network throughput while delivering superior data integrity.

The Intel® Ethernet CNA X520 family does it all: LAN, FCoE, and iSCSI; truly delivering on the promise of unified networking—on a single adapter.

## Reliable Performance

The X520 family of adapters includes a number of advanced features that enable it to provide industry-leading performance and reliability.

## Intelligent Offloads

With the use of intelligent offloads, Intel is able to target processing functions enabled with hardware assists or bypass offloads to optimize performance, scalability, and flexibility. The adapter's stateless offloads accelerate TCP/IP traffic while preserving processor computing power for critical applications. By doing operations in hardware, rather than local micro-processor, the X520-DA2 performance is able to scale with the processor generational gains.

## Security Optimizations

The adapters support IPsec offload for Microsoft's\* Network Access Protection (NAP), Active Directory,\* and future security capabilities in Windows 7\* and Windows 8\*. An X520 adapter enables customers to run a secure network environment without sacrificing performance.

## Designed for Multi-Core Processors

Support for technologies such as multiple queues, receive-side scaling, multiple MSI-X vectors, and Low Latency Interrupts enable the X520 family of adapters provide high-performance, 10 Gigabit connectivity in multi-core server platforms. These technologies distribute network processing across multiple CPU cores, improving overall performance.

## Intel® Ethernet Flow Director

Today's data centers depend on the multiprocessing and high performance capability of servers to increase system throughput, responsiveness, and reliability through the introduction of additional hardware threads, CPUs, or cores. But in a multiprocessing environment, it is essential to ensure a coordinated affinity of protocol processing and network applications on the same target cores. This affinity significantly reduces contention for shared resources, minimizes software synchronization overheads between cores, and enhances cache efficiency.

Receive Side Scaling (RSS) resolves the single-processor bottleneck by allowing the receive side network load from a network adapter to be shared across multiple cores. RSS enables packet receive-processing to scale with the number of available cores. However, RSS has a limitation—it cannot steer an incoming network flow to the same core on which its application process resides. RSS does not maintain the Traffic Flow → Core (Application) relationship. If an application is running on one core, while RSS has scheduled receive traffic to be on another core, poor cache efficiency and significant core-to-core synchronization overheads will result. The overall system performance can be significantly degraded.

Intel® Ethernet Flow Director, and the Application Target Routing (ATR) service found in Intel Ethernet Controllers, is an advanced network offload technology that provides the benefits of parallel receive processing in multiprocessing environments and automatically steers incoming network data to the same core on which its application process resides. Intel Ethernet Flow Director and ATR preserve the Traffic Flow → Core (Application) relationship. As a result, Intel Ethernet Flow Director and ATR can significantly lower latency and improve CPU utilization.

Intel Ethernet Flow Director allows administrators to define "signature filters" for the ATR service on the Intel Ethernet Controller to use to ensure that all packets in a TCP flow are processed by a single core. This intelligent offload capability supports advanced filters that direct receive packets by their flows to different queues and enables tight control on routing a flow in the platform. It matches flows and CPU cores for flow affinity and supports multiple parameters for flexible flow classification and load balancing.

For today's demanding virtualized data center environments, the new X520-DA2 adapter family delivers ultimate flexibility and scalability.

General Features	Benefits
Intel® 82599 Ethernet Controller	<ul style="list-style-type: none"> <li>Industry-leading, energy-efficient design for next-generation 10 Gigabit performance and multi-core processors</li> </ul>
SFP+ Connectivity	<ul style="list-style-type: none"> <li>X520 Adapters with SFP+ connections support 10GBASE-SR, 10GBASE-LR and SFP+ Copper Direct Attach physical media</li> </ul>
iSCSI remote boot support	<ul style="list-style-type: none"> <li>Provides centralized storage area network (SAN) management at a lower cost than other iSCSI solutions</li> </ul>
Fibre Channel over Ethernet (FCoE) Support	<ul style="list-style-type: none"> <li>FCoE Boot and Data Center Bridging ready. Contact your IBM representative for enabling.</li> </ul>
Support for most network operating systems (NOS)	<ul style="list-style-type: none"> <li>Enables widespread deployment</li> </ul>
RoHS-compliant	<ul style="list-style-type: none"> <li>Complies with the European Union directive 2002/95/EC to reduce the use of hazardous materials</li> </ul>
Intel® PROSet Utility for Windows* Device Manager	<ul style="list-style-type: none"> <li>Provides point-and-click management of individual adapters, advanced adapter features, connection teaming, and virtual local area network (VLAN) configuration</li> </ul>
Time Sync (IEEE 1588*, 802.1as)	<ul style="list-style-type: none"> <li>Enables networked Ethernet equipment to synchronize internal clocks according to a network master clock; endpoint can then acquire an accurate estimate of the master time by compensating for link latency.</li> </ul>
Load balancing on multiple CPUs	<ul style="list-style-type: none"> <li>Increases performance on multi-processor systems by efficiently balancing network loads across CPU cores when used with Receive-Side Scaling (RSS) from Microsoft or Scalable I/O on Linux*</li> </ul>
IBM Embedded Adapter form factor	<ul style="list-style-type: none"> <li>Enables port expansion without using a PCI expansion slot</li> </ul>
I/O Features for Multi-Core Processor Servers	
Intel® Data Direct I/O (Intel® DDIO)	<ul style="list-style-type: none"> <li>Reduces memory accesses from I/O on local socket</li> <li>Speeds up CPU data transfer</li> <li>Accelerates inbound &amp; outbound data flows</li> </ul>
Intel® Ethernet Flow Director	<ul style="list-style-type: none"> <li>Intel Ethernet Flow Director and ATR can significantly lower latency and improve CPU utilization by preserving the affinity between the flow and the core where the application resides</li> </ul>
RSS—Receive Side Scaling	<ul style="list-style-type: none"> <li>Uses multiple queues for receive traffic</li> </ul>
Intel® Direct Cache Access (DCA)	<ul style="list-style-type: none"> <li>Enables the adapter to pre-fetch the data from memory, avoiding cache misses and improving application response time</li> </ul>
MSI-X support	<ul style="list-style-type: none"> <li>Minimizes the overhead of interrupts</li> <li>Load-balancing of interrupt handling between multiple cores/CPU's</li> </ul>
Low Latency Interrupts (LLI)	<ul style="list-style-type: none"> <li>Based on the sensitivity of the incoming data, the adapter can bypass the automatic moderation of time intervals between the interrupts</li> </ul>
Header splits and replication in receive	<ul style="list-style-type: none"> <li>Helps the driver focus on the relevant part of the packet without the need to parse it</li> </ul>
Multiple Queues: 128 Tx and Rx queues per port	<ul style="list-style-type: none"> <li>Network packet handling without waiting or buffer overflow providing efficient packet prioritization</li> </ul>
Tx/Rx IP, SCTP, TCP, and UDP checksum offloading (IPv4, IPv6) capabilities	<ul style="list-style-type: none"> <li>Lower processor usage</li> <li>Checksum and segmentation capability extended to new standard packet type</li> </ul>
TxTCP segmentation offload (IPv4, IPv6)	<ul style="list-style-type: none"> <li>Increased throughput and lower processor usage</li> </ul>
Receive and Transmit Side Scaling for the Windows environment and Scalable I/O for Linux* environments (IPv4, IPv6, TCP/UDP)	<ul style="list-style-type: none"> <li>This technology enables the direction of the interrupts to the processor cores in order to improve the CPU usage rate</li> </ul>
IPsec offload	<ul style="list-style-type: none"> <li>Offloads IPsec capability onto the adapter instead of the software to significantly improve I/O throughput and CPU usage (for Windows* 2008 Server, Vista*, Windows* 2008 Server R2, and Windows* 7)</li> </ul>
MACSec	<ul style="list-style-type: none"> <li>IEEE spec: 802.1ae</li> <li>Layer 2 data protection that provides encryption and authentication ability between two individual devices (routers, switches, etc.)</li> <li>MACSec is designed into the network adapter hardware. These adapters are prepared to provide MACSec functionality when the ecosystem is ready to support this new technology</li> </ul>
Interrupt Throttle Rate (ITR)	<ul style="list-style-type: none"> <li>ITR parameter controls how many interrupts each interrupt vector can generate per second.</li> </ul>
Jumbo frames	<ul style="list-style-type: none"> <li>Supports jumbo frames larger than the default of 1500 bytes</li> </ul>

## I/O Features for Multi-Core Processor Servers *continued*

Large Receive Offload (LRO)	<ul style="list-style-type: none"><li>Combines multiple Ethernet frames into a single receive in the stack, thereby potentially decreasing CPU utilization for receives</li></ul>
MAC and VLAN anti-spoofing	<ul style="list-style-type: none"><li>If a malicious driver attempts to send a spoofed packet, it is dropped by the hardware and not transmitted. An interrupt is sent to the PF driver notifying it of the spoof attempt.</li></ul>
Flow Control	<ul style="list-style-type: none"><li>Ethernet Flow Control (IEEE 802.3x) support for capable link partner</li></ul>
HW-Based Receive Side Coalescing (RSC)	<ul style="list-style-type: none"><li>Merges multiple frames from the same IPv4 TCP/IP flow into a single structure that can span one or more descriptors</li></ul>

## Virtualization Features

VMDq	<ul style="list-style-type: none"><li>Offloads the data-sorting functionality from the Hypervisor to the network silicon, improving data throughput and CPU usage.</li><li>Provides QoS feature on the Tx data by providing round-robin servicing and preventing head-of-line blocking</li><li>Sorting based on MAC addresses and VLAN tags</li></ul>
PC-SIG SR-IOV implementation (up to 64 virtual functions per port)	<ul style="list-style-type: none"><li>Provides an implementation of the PCI-SIG standard for I/O Virtualization. The physical configuration of each port is divided into multiple virtual ports. Each virtual port is assigned to an individual virtual machine directly by bypassing the virtual switch in the Hypervisor, resulting in near-native performance.</li><li>Integrated with Intel® VT for Directed I/O (Intel® VT-d) to provide data protection between virtual machines by assigning separate physical addresses in the memory to each virtual machine.</li></ul>
Flexible Port Partitioning, 64 Virtual Functions per port	<ul style="list-style-type: none"><li>Virtual Functions (VFs) appear as Ethernet Controllers in Linux OSes that can be assigned to VMs, Kernel processes or teamed using the Linux* Bonding Drivers</li></ul>
Virtual Machine Load Balancing (VLMB)	<ul style="list-style-type: none"><li>Virtual Machines Load Balancing (VMLB) provides traffic load balancing (Tx and Rx) across Virtual Machines bound to the team interface, as well as fault tolerance in the event of switch, port, cable, or adapter failure</li></ul>
Advanced Packet Filtering	<ul style="list-style-type: none"><li>24 exact-matched packets (unicast or multicast)</li><li>4096-bit hash filter for unicast and multicast frames</li><li>Lower processor usage</li><li>Promiscuous (unicast and multicast) transfer mode support</li><li>Optional filtering of invalid frames</li></ul>
VLAN support with VLAN tag insertion, stripping and packet filtering for up to 4096 VLAN tags	<ul style="list-style-type: none"><li>Ability to create multiple VLAN segments</li></ul>

## Manageability Features

PreBoot eXecution Environment (PXE) Support	<ul style="list-style-type: none"><li>Enables system boot up via the LAN (32-bit and 64-bit)</li><li>Flash interface for PXE image</li></ul>
Simple Network Management Protocol (SNMP) and Remote Network Monitoring (RMON) statistic counters	<ul style="list-style-type: none"><li>Easy system monitoring with industry-standard consoles</li></ul>
iSCSI and FCoE Boot	<ul style="list-style-type: none"><li>Enables system boot-up via iSCSI or FCoE</li><li>Provides additional network management capability</li></ul>
Watchdog timer	<ul style="list-style-type: none"><li>Gives an indication to the manageability firmware or external devices that the chip or the driver is not functioning</li></ul>

## Specifications

### General

Connectors	<ul style="list-style-type: none"><li>Two SFP+ Cages Supporting 10GBASE-SR and 10GBASE-LR Transceivers</li><li>SR and LR Optic Transceivers have LC fiber-optic connectors</li><li>SFP+ Direct Attach cables (E10G42BTDA)</li></ul>
Network Standards	<ul style="list-style-type: none"><li>IEEE 802.3:<ul style="list-style-type: none"><li>–10GBASE-SR</li><li>–10GBASE-LR</li></ul></li><li>SFF-8431:<ul style="list-style-type: none"><li>–110GSFP+Cu (a.k.a. Direct Attach)</li></ul></li></ul>

## Adapter Product Features

Intel® PROSet Utility	For easy configuration and management
Plug and play specification support	Standard
Receive Side Scaling	Multiple Rx queues enable the efficient distribution of network receive processing across multiple CPUs in multiprocessor system
Direct Cache Access (DCA)	The I/O device activates a pre-fetch engine in the CPU that loads the data into the CPU cache ahead of time, before use, eliminating cache misses and reducing CPU load

## Advanced Software Features

Adapter Fault Tolerance (AFT)
Switch Fault Tolerance (SFT)
Adaptive Load Balancing (ALB)
Teaming support
IEEE 802.3ad (link aggregation control protocol)
PCIe Hot Plug*/Active Peripheral Component Interconnect (PCI)
IEEE 802.1q* VLANs
IEEE 802.3 2005* flow control support
Tx/Rx IP, TCP, & UDP checksum offloading (IPv4, IPv6) capabilities (Transmission Control Protocol (TCP), User Datagram Protocol (UDP), Internet Protocol (IP))
IEEE 802.1p*
TCP segmentation/large send offload
MSI-X supports Multiple Independent Queues
Interrupt moderation
IPv6 offloading—checksum and segmentation capability extended to new standard packet type

## Technical Features

Data rate supported per port	<ul style="list-style-type: none"> <li>▪ Optical: 1 GbE/10 GbE</li> <li>▪ Direct Attach: 10 GbE</li> </ul>
Bus type	PCI Express 2.0 (5.0 GT/s)
Bus width	4-lane PCI Express and 8-lane PCI Express
Interrupt levels	INTA, MSI, MSI-X
Hardware certifications	FCC B, UL, CE, VCCI, BSMI, CTICK, KCC
Controller-processor	Intel® 82599 Ethernet Controller
LED Indicators	LINK (solid) and ACTIVITY (blinking) LINK SPEED (green=10 Gbps; yellow=1 Gbps)

Power Consumption	Maximum	Typical
Dual-port 10GBASE-SR	7.547W	6.974W
Dual-port 1000BASE-SX	5.700W	5.033W
Dual-port 10GBASE-LR	8.114W	7.061W
Dual-port 1000BASE-LX	6.377W	5.392W
Dual-port 10GBASE-SFP+	6.551W	5.553W

## Environment

Operating temperature	0 °C to 55 °C (32 °F to 131 °F)
Air flow	Minimum of 100 LFM required
Storage temperature	-40 °C to 70 °C (-40 °F to 158 °F)
Storage humidity	Maximum: 90% non-condensing relative humidity at 35 °C

## Physical

Dimensions	6.6 in. x 2.7 in. x 0.6in. (168 mm x 69mm x 15mm)
------------	---

## Network Operating System (NOS) Support

Operating System	IA-32	x86-64	IA-64
Windows 7	X	X	N/A
Windows Server 2008 R2	N/A	X	X
Windows Server 2008 R2 Core	N/A	X	X
Linux SLES 11 SP1	X	X	X
Windows Vista* SP2	X	X	N/A
Windows Server* 2003 SP2	X	X	X
Windows Server 2008 SP2	X	X	X
Windows Server 2008 SP2 Core	X	X	N/A
Linux* Stable Kernel version 2.6	X	X	X
Linux RHEL 4.8	X	X	N/A
Linux RHEL 5.5	X	X	X
Linux SLES 10 SP3	X	X	X
Linux SLES 11	X	X	X
FreeBSD* 8.0	X	X	X
EFI* 1.1	N/A	N/A	X
UEFI* 2.1	N/A	X	X

### Intel® Network Adapter X520 Product Codes

Configuration	Product Code
X520-DA2	49Y7981 <sup>1</sup>

### Intel® Ethernet SFP+ Twinaxial Cable Product Codes

Cable Length (m)	Product Code
1	XDACBL1M
2	XDACBL3M
3	XDACBL5M

### Intel® Ethernet SFP+ Optic Product Codes

Configuration	Product Code
SR Optic	E10GSFPSR
LR Optic	E10GSFPLR

#### IBM Systems supporting the Intel Ethernet X520-DA2 Embedded Converged Network Adapter:

System	Intel Ethernet X520-DA2 Embedded CNA
IBM 3550 M4 (7914)	Yes
IBM 3650 M4 (7915)	Yes
dx360 M4 (7912)	Yes

## For Product Information

To speak to a customer service representative, please call 1-800-538-3373 (U.S. and Canada) or visit [support.intel.com/support/go/network/contact.htm](http://support.intel.com/support/go/network/contact.htm) for the telephone number in your area.

For additional product information on Intel Networking Connectivity products, visit [www.intel.com/go/ethernet](http://www.intel.com/go/ethernet).

## Customer Support

Intel® Customer Support Services offers a broad selection of programs including phone support and warranty service. For more information, contact us at [support.intel.com/support/go/network/adapter/home.htm](http://support.intel.com/support/go/network/adapter/home.htm).

(Service and availability may vary by country.)

For more information, visit [www.IntelEthernet-IBM.com](http://www.IntelEthernet-IBM.com) or contact your IBM sales representative.

<sup>1</sup> Requires a system that supports VMDq.

The information contained in this document, including all instructions, cautions, and regulatory approvals and certifications, is provided by Intel and has not been independently verified or tested by IBM. IBM cannot be responsible for damage caused as a result of either following or failing to follow these instructions. All statements or claims regarding the properties, capabilities, speeds or qualifications of the part referenced in this document are made by Intel and not by IBM. IBM specifically disclaims knowledge of the accuracy, completeness or substantiation for any such statements. All questions or comments relating to such statements or claims should be directed to Intel. Visit [www.ibm.com](http://www.ibm.com) for more information.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order. Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web Site at <http://www.intel.com/design/literature.htm>.

Copyright © 2013 Intel Corporation. All rights reserved. Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and other countries. Printed in USA

\*Other names and brands may be claimed as the property of others.

0513/MBR/RCE/PDF



Third-party information brought to you courtesy of IBM.

