


White Paper

3G Converged-NICs A Platform for Server I/O to Converged Networks

This document helps those responsible for connecting servers to networks achieve network convergence by providing an overview of key industry dynamics, Ethernet capabilities with data center bridging, and an understanding of how C-NICs will emerge as the platform for connectivity to converged networks.



June 2010



Data Center Infrastructure Consolidation

Infrastructure consolidation is a key data center initiative because it simplifies operations and saves money. That's why IT organizations routinely standardize infrastructure to consolidate training, deployment, and management. This practice extends to host connectivity where standardizing on one type of hardware platform and one driver image is common for LAN, SAN, and high-performance computing (HPC) network adapters. In a recent survey of IT professionals with hands-on experience with host adapters, 38% said that all of their servers were configured with a standard 1-GbE NIC and driver image. Another 18% said that up to 75% of their servers were configured with a standard 1-GbE NIC image. Another 18% said that up to 75% of their servers were configured with a standard 1-GbE NIC image. Another 18% said that up to 75% of their servers were configured with a standard 1-GbE NIC image.

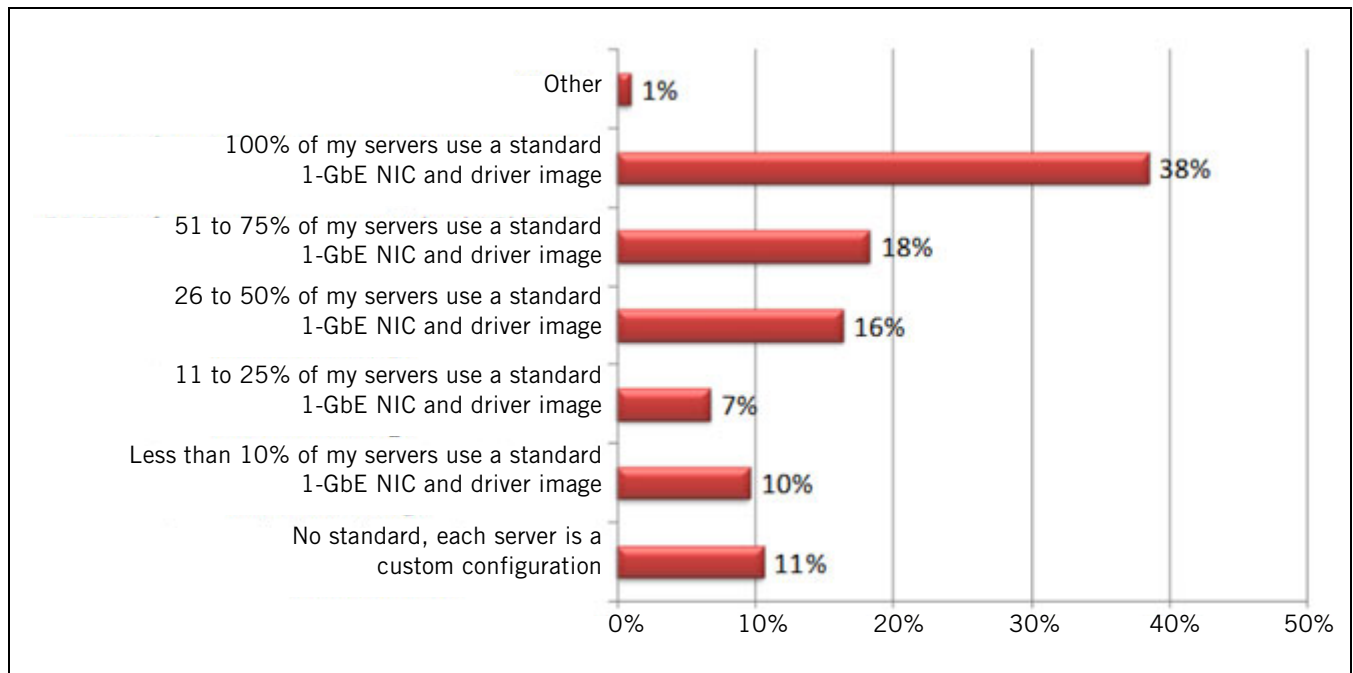


Figure 1: IT Pro Survey Results

“What percentage of your servers are configured with a standard 1-GbE NIC and driver image?”

Server I/O Consolidation Has Not Been Possible

Today, dramatically different classes of network controller ASICs are required to support the specific needs of Ethernet LANs, Ethernet/iSCSI SANs, Fibre Channel SANs, and InfiniBand® HPC clusters. I/O consolidation has not been possible because prior generations of network controllers were not designed for the variety of network types that would allow universal deployments within the LAN-on-Motherboard (LOM) or adapter card system environments.

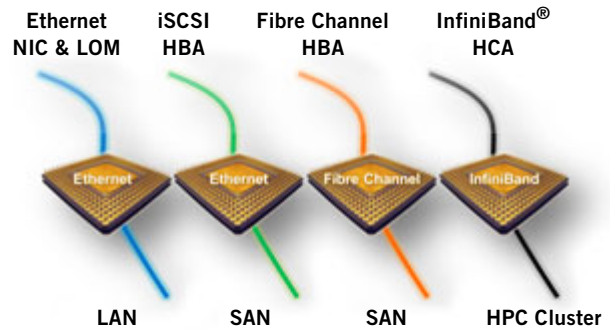


Table 1: Types of I/O Adapters for Networking in the Data Center

| Network I/O Types | Network | Unique Capability |
|-------------------------------------------------|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LOM (LAN-on-Motherboard) | LAN | 1/10 GbE controllers embedded in most servers for communications with clients, remote server management, file-oriented network attached storage (NAS), and block-oriented iSCSI storage |
| Ethernet NIC (Network Interface Card) | LAN | 1/10 GbE PCI-E and blade server mezzanine cards used for communications with clients, remote server management, and file-oriented NAS |
| iSCSI HBA (Host Bus Adapter) | SAN | 1/10 GbE PCI-E and blade server mezzanine cards used for communications with block-oriented iSCSI storage |
| Fibre Channel HBA (Host Bus Adapter) | SAN | 4/8 GbE PCI-E and blade server mezzanine cards used for communications with block-oriented Fibre Channel storage |
| InfiniBand HCA (Host Channel Adapter) | HPC Cluster | 20/40 GbE PCI-E and blade server mezzanine cards used for low-latency inter-processor communications (IPC) required in high-performance computing clusters |

A New Generation of 10-GbE Technology


Products based on the first generation of 10-GbE controller ASICs have been commercially available since 1994. A new class of Ethernet controller ASICs is emerging in 2010 to support the second generation of 10-GbE technology. Key components of the new 10 Gb Ethernet are data center bridging, iSCSI offload, Fibre Channel over Ethernet, Remote Direct Memory Access (RDMA), and 10GBASE-T.

Data Center Bridging

iSCSI for iSCSI SANs

FCoE for FC SANs

iWarp for HPC Clusters



Data Center Bridging

Ethernet is designed to be a best-effort network that may drop packets or deliver packets out of order when the network or devices are busy. Extensions to the Ethernet protocol provide reliability without incurring the penalties of TCP. Additionally, they provide more control of bandwidth allocation and, thus, a more effective use of bandwidth. These enhancements make Ethernet C-NICs viable host adapters for Fibre Channel SAN and HPC cluster traffic.

Table 2: Ethernet Enhancements That Comprise Data Center Bridging

| IEEE Protocol | What It Does |
|------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| Priority Flow Control (PFC) P802.1Qbb | A link-level priority flow control mechanism to pause different classes of traffic |
| Enhanced Transmission Selection (ETS) P802.1Qaz | Manages the bandwidth between different traffic classes for multiprotocol links |
| Congestion Management (CM) P802.1Qau | Provides end-to-end congestion management for the Fibre Channel over Ethernet (FCoE) protocol |
| Data Center Bridging Capabilities Exchange Protocol (DCBCXP) 802.1Qaz | A discovery and capabilities exchange protocol to ensure consistent configuration of enhanced Ethernet devices across a network |

iSCSI HBA Offload

A server can support connectivity to an iSCSI SAN by configuring a NIC with a software initiator that runs on the server. A C-NIC with iSCSI HBA offload removes the burden of iSCSI protocol processing from the server and frees the processor resources to support the business applications for which the server was purchased.

Fibre Channel over Ethernet

This protocol was developed so that IT organizations could leverage their Ethernet training and equipment to support connectivity to Fibre Channel storage. This is accomplished with the FCoE protocol that encapsulates Fibre Channel frames over Ethernet networks. C-NICs that support FCoE connect to Fibre Channel storage through an Ethernet switch with FCoE bridging or directly to native FCoE storage through a traditional Ethernet switch.

10GBASE-T

In the near future, most data centers will migrate from 1-GbE networks to 10-GbE networks. A mass migration will be accelerated with 10GBASE-T technology that enables the use of familiar and low-cost RJ45 connectors. C-NICs with RJ45 connectors and 10GBASE-T allow data center managers to deploy 10-GbE networks and storage using CAT6A copper cables already present in most of the world's data centers. Without 10GBASE-T, miles of new optical cables and/or more expensive optical connectors must be run to migrate an entire data center to 10 Gb speeds. A recent end-user survey by IT Brand Pulse indicates that IT professionals understand the benefits of 10GBASE-T technology. Approximately 50% said they would prefer to connect their 10-GbE adapters with 10GBASE-T.

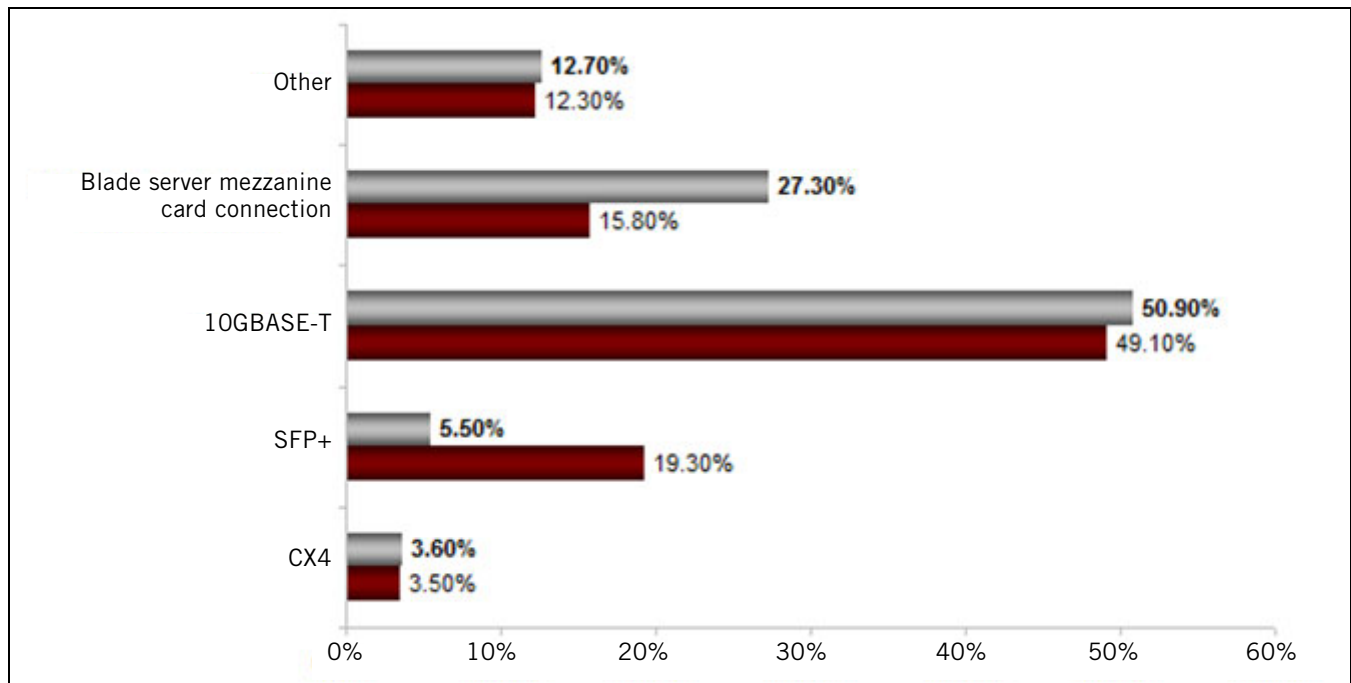


Figure 2: IT Pro Survey Results—Dell and HP End Users

“To which 10-GbE adapters do you prefer to connect?”

Convergence on Ethernet

Eventually the new generation of Ethernet, viable for Fibre Channel SANs and HPC clusters, will be pervasive in the data center. During convergence, islands of Fibre Channel, Ethernet, and InfiniBand adapters, switches, and storage will converge onto Ethernet.

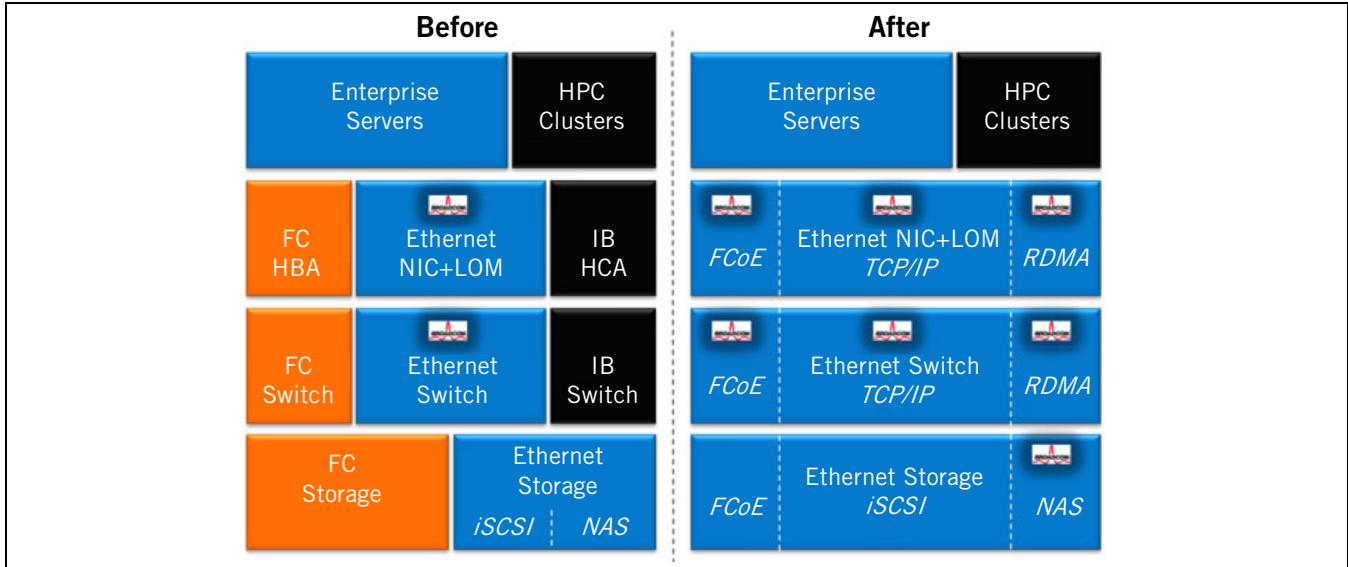


Figure 3: Data Center Convergence on Ethernet

The process of convergence on Ethernet is a continuum consisting of many milestones achieved by vendors and IT organizations over the course of several years. Convergence on Ethernet is already being implemented by organizations deploying LANs and iSCSI SANs.

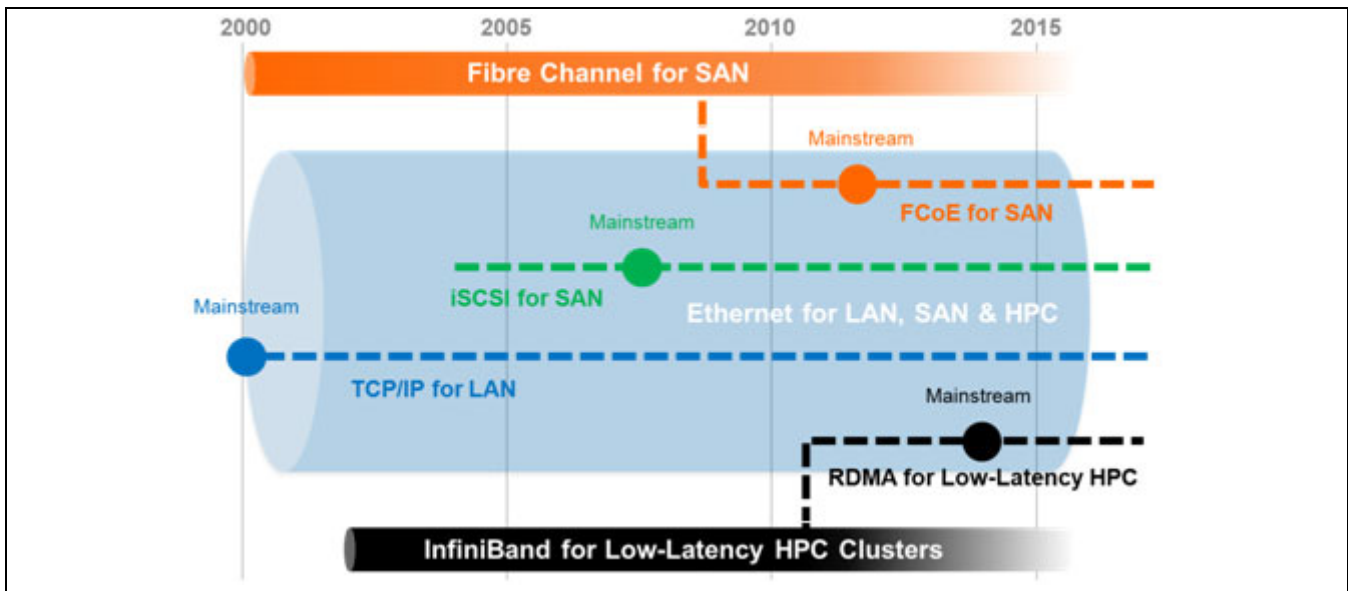


Figure 4: Timeline for Convergence on Ethernet

Server Virtualization Driving the Need for 10 GbE

Server virtualization has clearly moved to the core of data center architectures and is driving the emergence of dense computing nodes consisting of multiprocessor and multicore servers loaded with virtual machines. As a result, data center managers are dealing with a massive aggregation of I/O and a corresponding need for more network bandwidth. The 8x growth in 10-GbE NIC port shipments expected between 2009 and 2013 indicates their response will be to dramatically increase the deployment of 10-GbE networks.

Today most servers running server virtualization software are configured with multiple 1 GbE NIC ports to support the increased I/O load from multiple virtual machines. A recent focus group said their virtualized servers were configured with an average of eight 1 GbE NIC ports. In a related survey, data center managers overwhelmingly selected server virtualization as the primary driver for 10 GbE adoption. Members of the focus group expressed a strong interest in a 10 GbE C-NIC platform to provide additional network bandwidth and to reduce the number of ports for reduced power and management complexity.

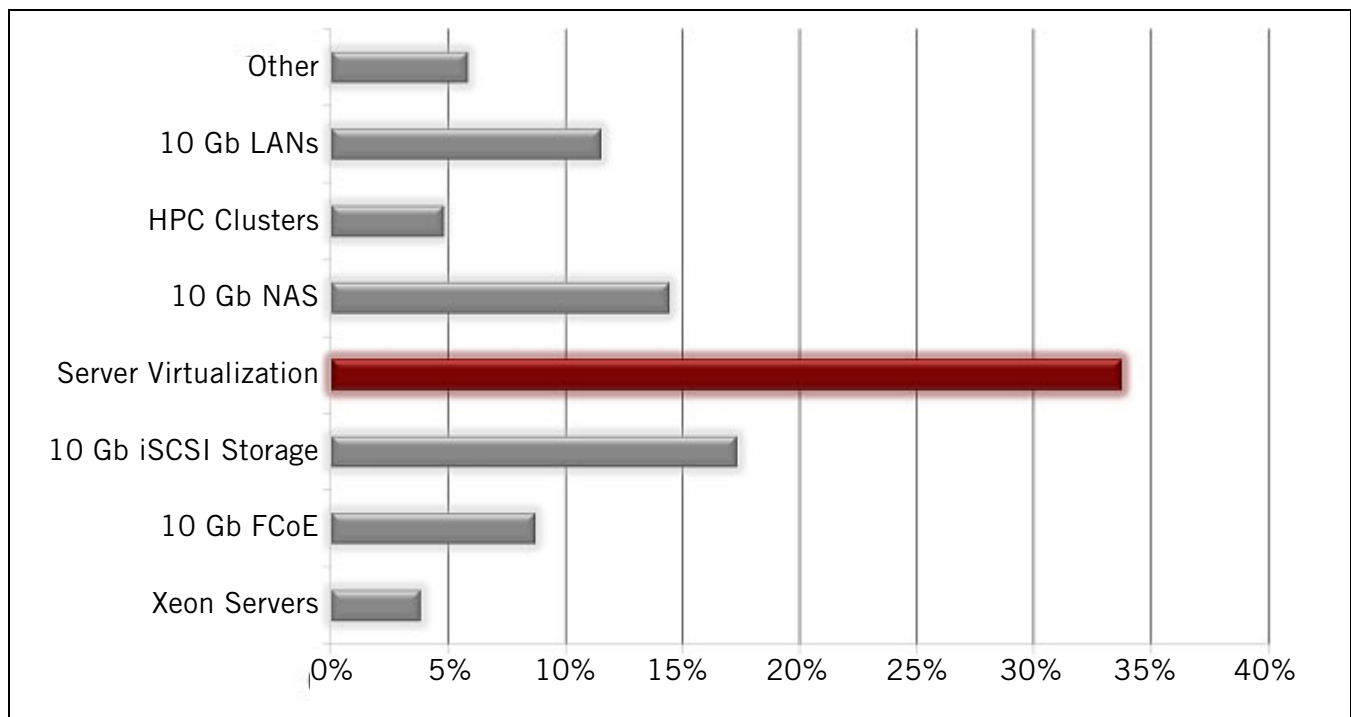


Figure 5: IT Pro Survey Results—Server Virtualization

“Which data center application will drive your adoption of 10 GbE?”



10 GbE for Server I/O to Converged Networks

An entire networking industry is at the threshold of a massive migration from disparate LAN, SAN, and HPC networks to converged networks based on Ethernet with data center bridging. One result is that multiple types of host adapters will be replaced by third-generation converged NICs that offer a combination of the following five features and benefits found only in 3G C-NICS:

1. Hardened LAN Connectivity

3G C-NICs and LOM make up 88% of the network I/O ports in the data center today.




Figure 6: Data Center Network Adapter Types

2. Backward Compatibility

3G C-NICs and LOM deliver backwards compatibility with 1-GbE adapters including training, drivers, installation scripts, and management software.

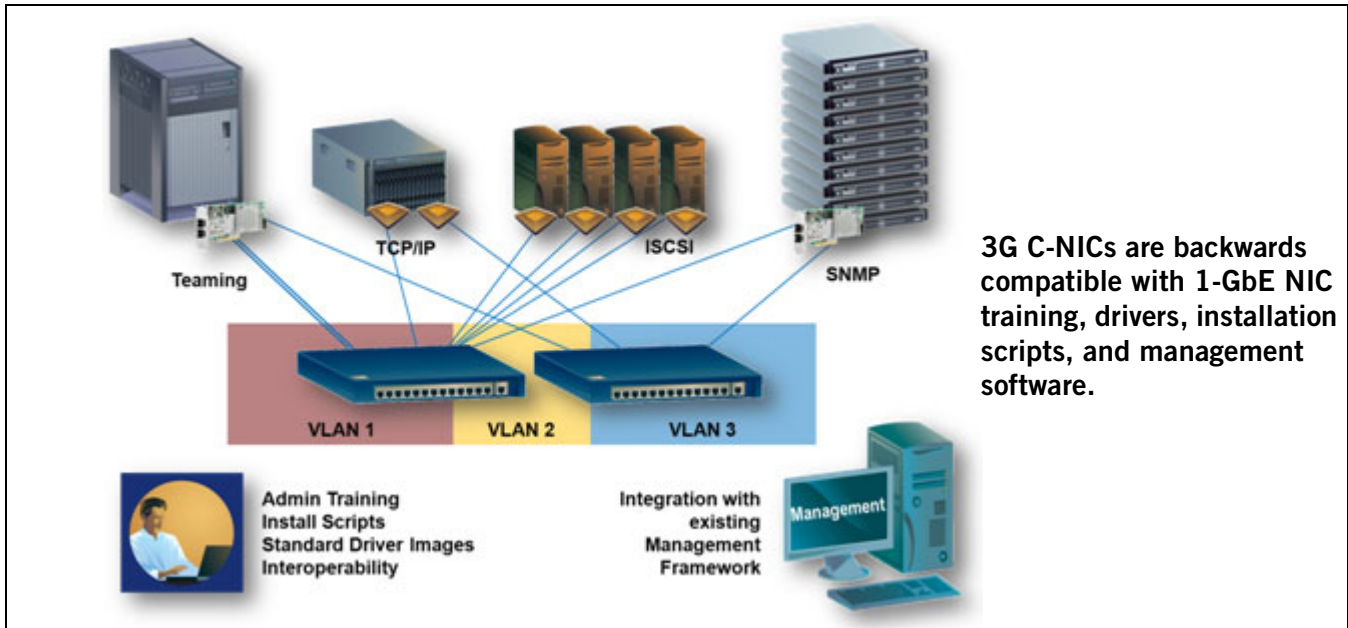


Figure 7: 3G C-NIC Backwards Compatibility

3. Hardened for iSCSI SANs

In 2010 iSCSI storage arrays will comprise about 38% of all storage arrays. By 2013 about half of all storage arrays will be iSCSI. 3G C-NICs and LOM are hardened for fast iSCSI SAN environments with a 1/10 GbE adapter run-rate of over 1M iSCSI offload HBA ports per year in 2010...and growing rapidly.

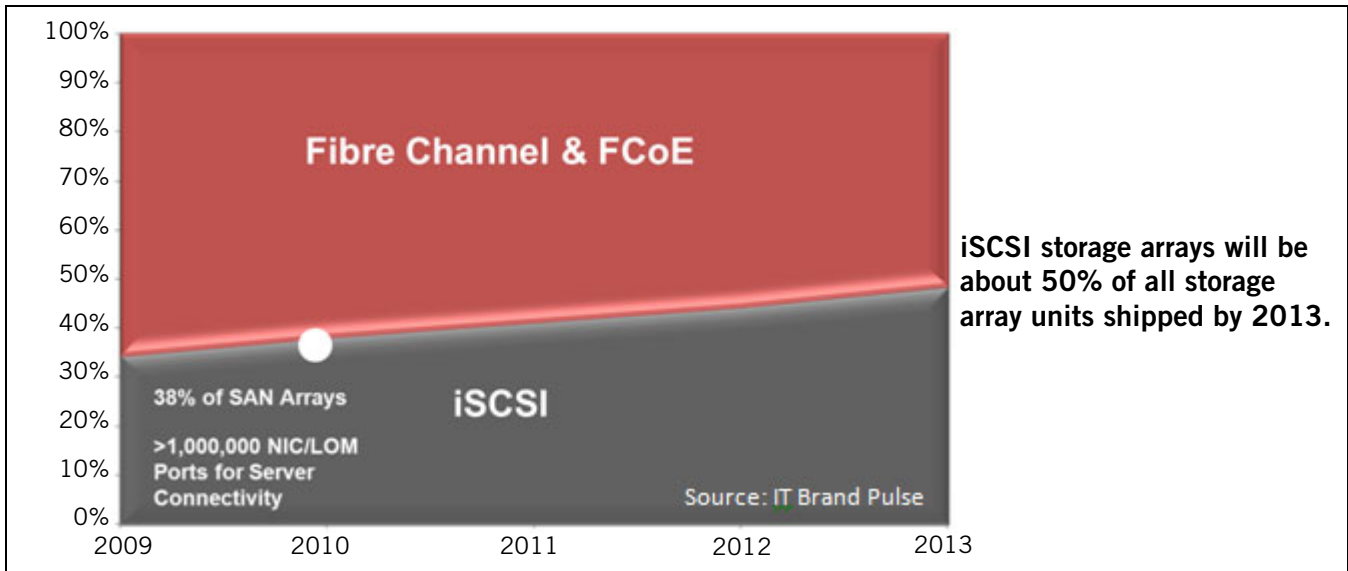


Figure 8: External Storage Array Unit Shipments

4. Bridge to Fibre Channel SANs

3G C-NICs and LOM offer simultaneous connectivity to NAS, iSCSI SANs, *and* Fibre Channel SANs.

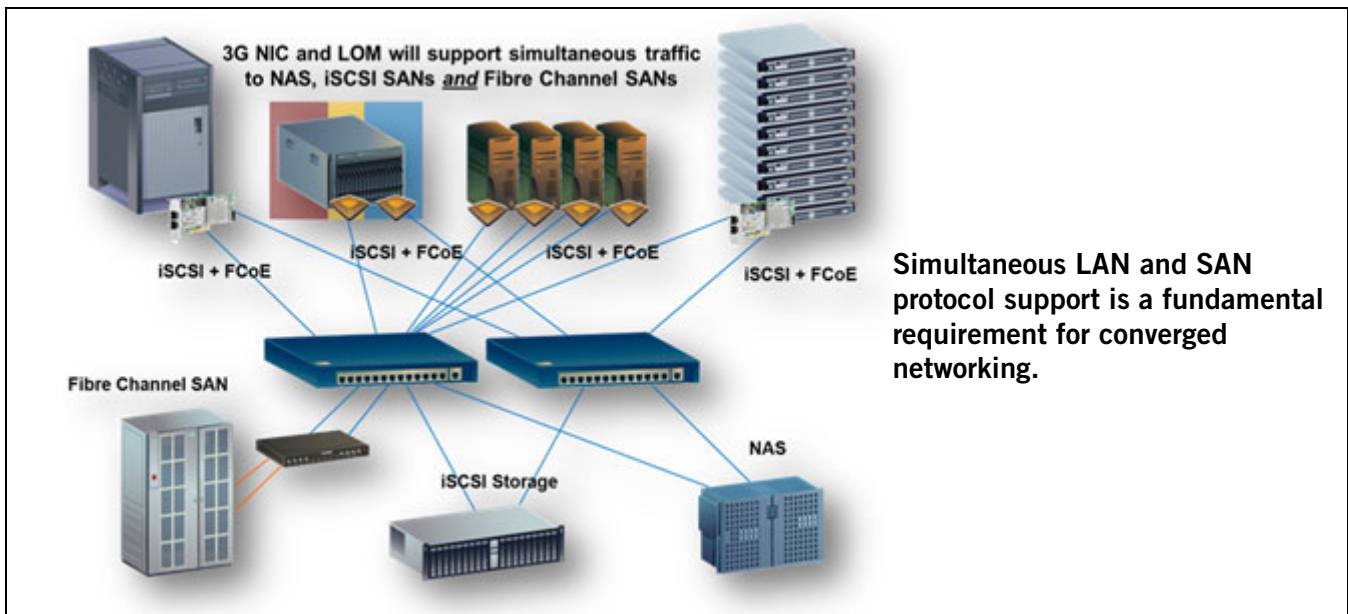


Figure 9: Simultaneous LAN and SAN Protocol Support

5. Forward Looking

3G C-NICs and LOM are backwards compatible with 1 GbE networks and forward looking with an architecture designed to provide lower power and lower cost solutions to data center managers.

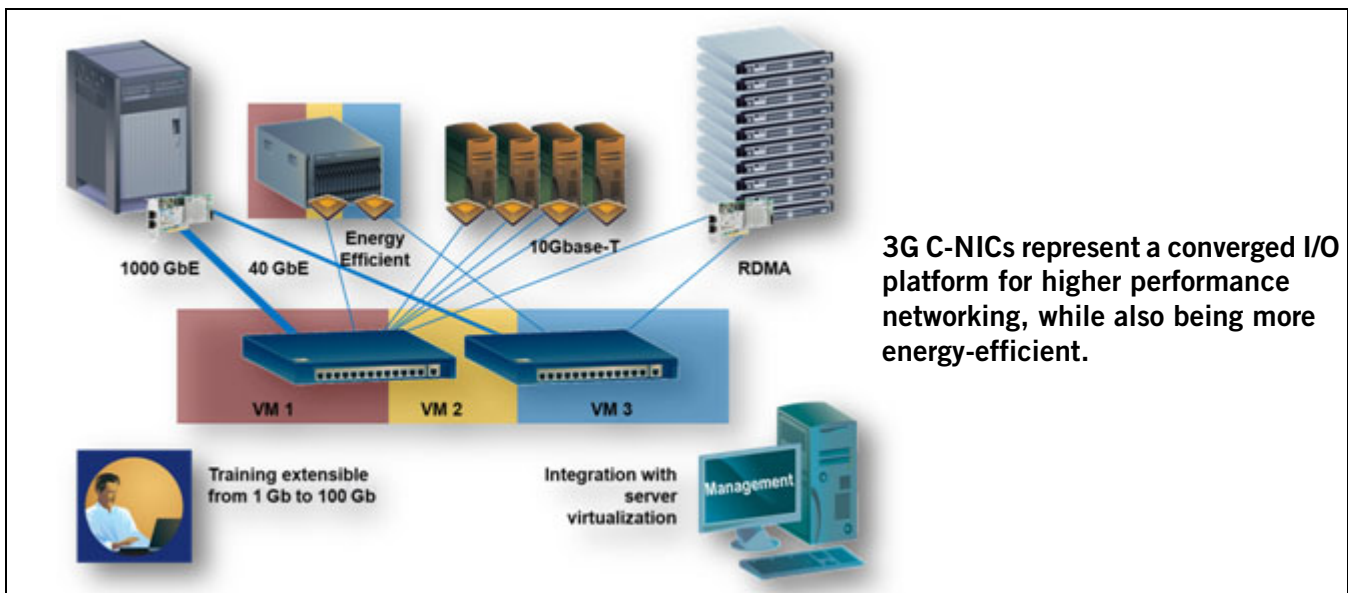


Figure 10: 3G C-NIC Converged I/O Platform



A No-Compromise Converged I/O Platform Checklist

The following checklist identifies the features needed for a 10 GbE C-NIC to serve as a platform for I/O convergence in your data center:



Data Center Bridging (DCB) support for all protocols: DCB is what enables the *Lossless Ethernet* capability needed for storage networking. Some NICs and Converged Network Adapters (CNAs) support DCB exclusively on one protocol at a time. 3G C-NICs are designed to support DCB while running all SAN, LAN, and offload protocols simultaneously.



Simultaneous LAN and SAN protocol support: For data centers that deploy LANs and a tiered SAN model that includes iSCSI and FC storage, a converged NIC must be able to run the TCP/IP, iSCSI, and FCoE protocols simultaneously.



Hardened LAN and iSCSI stack: Since 88% of data center ports are LAN ports, it's best to converge on a 10 GbE C-NIC platform hardened from millions of installed 1 GbE NICs.



Common drivers for 1 GbE and 10 GbE: It's possible to experience a seamless migration from 1 GbE to 10 GbE by using C-NICs that share a common driver image.



Full offload protocol support to reduce CPU utilization: Some NICs and CNAs support iSCSI software initiators, but only 3G C-NICs provide full-offload engines to reduce CPU utilization and maximize CPU usage for applications. Offload protocols can reduce CPU utilization by 70%, which also means lower power draw by the server.



High Performance, coupled with low CPU utilization: Some NICs and CNAs provide high Input/Output Operations Per Second (IOPS), while some NICs and CNAs minimize CPU utilization. Only 3G C-NICs do both, allowing a network to perform without sacrificing valuable CPU cycles allocated for applications and virtualization.



Low power, Wake-On-LAN, and Energy Efficient Ethernet: Implementation of the IEEE 802.3az specification for Energy Efficient Ethernet (EEE) reduces server power usage when workloads are low. Adhering to Wake-on-LAN and EEE standards helps free-up data center budgets and is an important contributor to a greener data center.



10GBASE-T: CAT6A cabling is pervasive for traditional Ethernet networking. Only 3G C-NICs support 10GBASE-T connections, allowing data center managers to use existing CAT6A cables, longer distances, lower RJ45 connector costs, and existing tools.

Summary

If you're a server, storage, or network architect or administrator responsible for host connectivity to a LAN, SAN, or HPC cluster, now is a good time to begin planning your migration to converged 10 GbE networks and C-NICs.

More on Broadcom's 10 GbE

For more information about Broadcom's 10 GbE controllers and adapters visit:

<http://www.broadcom.com/products/Ethernet-Controllers/Enterprise-Server>

Related Articles

iSCSI HBA offload and 10GBASE-T:

http://www.infostor.com/index/blogs_new/Frank_Berry/blogs/infostor/frank-berry_s_blog/post987_855043080666472366.html

<http://www.networkcomputing.com/next-gen-network/broadcom-unveils-worlds-1st-converged-nic-with-iscsi-offload-on-10gbase-t.php>

Convergence and Virtualization:

http://www.infostor.com/index/articles/display/4348346471/articles/infostor/san/fibre-channel/2010/april-2010/fcoe-i_o_convergence.html

Broadcom®, the pulse logo, **Connecting everything®**, and the Connecting everything logo are among the trademarks of Broadcom Corporation and/or its affiliates in the United States, certain other countries and/or the EU. Any other trademarks or trade names mentioned are the property of their respective owners.

Connecting
everything®



BROADCOM CORPORATION

5300 California Avenue

Irvine, CA 92617

© 2010 by BROADCOM CORPORATION. All rights reserved.

Phone: 949-926-5000

Fax: 949-926-5203

E-mail: info@broadcom.com

Web: www.broadcom.com