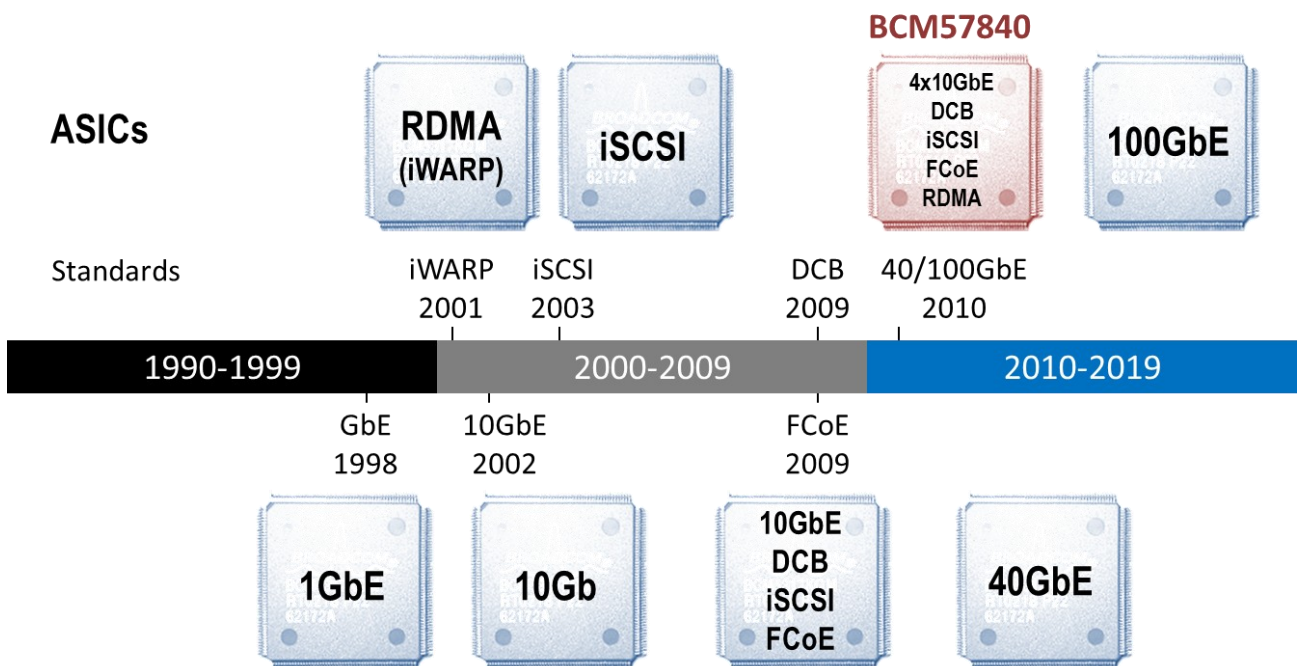


Broadcom Raises the Bar for CNA Technology

By Frank Berry

Broadcom Corporation recently unveiled the BCM578x0, its latest generation of 10GbE ASICs for LAN-on-Motherboard (LOM) and NIC applications, and the first in a new class of Converged Network Adapters (CNAs), a.k.a. Converged NICs (C-NICs). The new family of ASICs includes several innovations that raise the bar for CNA technology and give us a glimpse at the capabilities of the next generation of high performance Ethernet adapters.

Ethernet LOM/NIC ASIC—Industry Timeline



The standards are in place for a new generation of ASICs that will serve as the platform for complete network convergence. The Broadcom BCM578x0 is first in a new class of ASICs with at least 40Gb of bandwidth, and with hardware offload of TCP LAN, iSCSI SAN, FCoE SAN and RDMA HPC protocols.

5 Reasons Why a New LOM/NIC Platform is Needed

There are five key reasons why data center managers will need a new generation of LOM and NIC in their virtualized servers.


1. Servers with 8+ Cores Can Drive Multiple 10GbE Ports

The days when servers didn't have the power to drive 10GbE traffic are gone. The new generation of 8 and 12-core processors from Intel and AMD are designed to drive I/O to-and-from servers loaded with virtual machines. According to Intel, their new 8-core Xeon 7500 can deliver up to 20x the performance of servers with single core processors.

2. Servers are Set to Scale 10GbE Ports with PCIe 3.0

Another reason the capabilities of this new generation of LOM and NIC will be coveted in data centers, is because they're speed matched with the latest generation of PCIe 3.0 I/O interconnects, which will appear in new servers beginning in 2011. The bi-directional bandwidth of an x8 PCIe bus is 64Gb, which means a single bus can support up to six 10GbE, or one 40GbE adapters running at full speed.

Ethernet LOM/NIC Ports per PCIe 3.0 Bus



PCIe Architecture	Interconnect Bandwidth	Bi-Directional Bandwidth	Total Bandwidth for x8 Link	10GbE LOM/NIC Ports Supported	40GbE LOM/NIC Ports Supported
PCIe 1.x	2Gb/s	4Gb/s	4Gb/s	0	0
PCIe 2.0	4Gb/s	8Gb/s	32Gb/s	3	0
PCIe 3.0	8Gb/s	16Gb/s	64Gb/s	6	1


5 Reasons Why a New LOM/NIC Platform is Needed (cont.)

3. Virtualized Servers Require Full Offload

VMware's inbox support for 10GbE iSCSI offload in 2010 was an acknowledgement that CPU should be preserved for Virtual Machines, and virtualized servers require adapters with full offload. Best practices in the future will include deployment of 10GbE adapters that offload all network protocol and virtual switch processing, including TCP/IP, iSCSI, FCoE, RDMA and Virtual Ethernet Bridging (VEB).


Offloading Protocol Processing for Virtualized App Servers

Ample CPU Resources for Protocol Processing and VM Switching on Application Servers



Application Server

All CPU Resources Needed for Servers Heavily Loaded with VMs and Applications



Protocol Processing & VM switching off-loaded to LOM and C-NICs


Virtualized Application Server

5 Reasons Why a New LOM/NIC Platform is Needed

4. Virtualized Servers Require More System Memory

As the average number of virtual servers per physical server grows, so does the average amount of memory configured per physical server. The new generation of Ethernet LOM ASICs is critically important to server designers and admins, because it frees motherboard space for more system memory.

Consolidating LOM ASICs Frees Space for System Memory

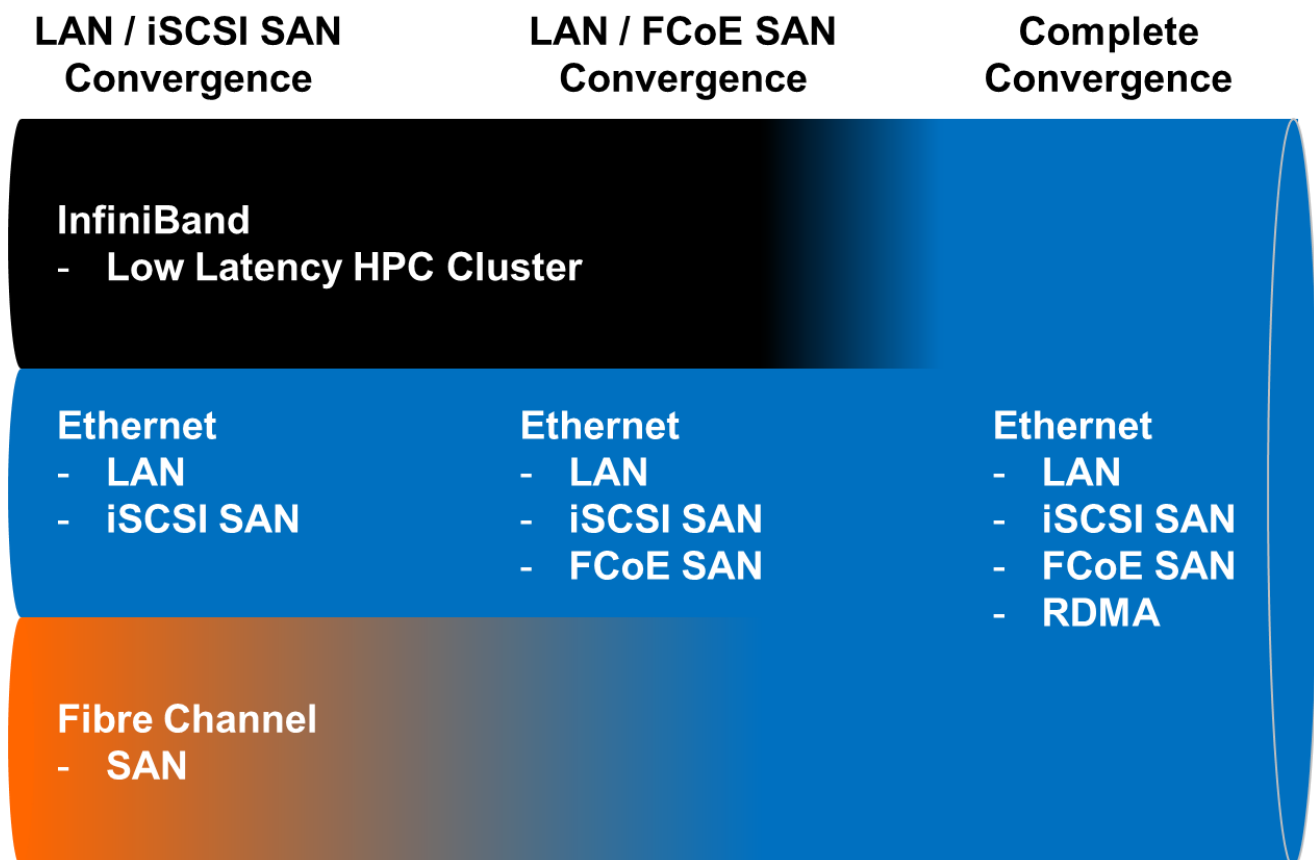


5 Reasons Why a New LOM/NIC Platform is Needed (cont.)

5. Complete Convergence on Ethernet Requires Lower Latency Ethernet

The current generation of 10GbE adapters and switches support Data Center Bridging (DCB) and FCoE for network convergence. However, complete network convergence requires support for the Remote Direct Memory Access (RDMA) protocol used for low-latency traffic found in high performance computing (HPC) clusters. When this final convergence building block is put into place, LANs, NAS, iSCSI SAN, FCoE SAN and HPC cluster traffic can converge onto Ethernet.

Support for RDMA Enables Complete Network Convergence



BCM57840: First in a New Class of LOM/NIC ASICs

The BCM578x0 responds to these five requirements, plus has the basic features needed for 40GbE adapter and storage ports in the near future.

Each chip in the BCM578x0 family features 40Gb of aggregate bandwidth, and support for the 64Gb x8 version of the new PCI Express (PCIe) 3.0 interconnect, which you can expect to see in servers later this year.

The space-saving quad-port version of the BCM578x0 also delivers twice the port density of competitive dual-port ASICs, in a footprint smaller than the smallest dual-port 10GbE ASICs from Brocade, Emulex, Intel and Mellanox.


Higher BCM578x0 “physical” port density allows 10GbE ports to scale cost effectively. Today, the deployment of virtual NIC technology is in its early stages of adoption and 5-10 1GbE ports are commonly configured on a midrange server. As server virtualization continues to drive the adoption of 10GbE over 1GbE, quad-port 10GbE ASICs will help server designers and admins cost-effectively scale up 10GbE LOM and NIC ports.

Higher “virtual” NIC port density will be a priority in the future. It is not uncommon for server admins to configure 20+ Virtual Machines per physical server. These configurations may need 20+ Ethernet ports. The BCM578x0 class of ASIC provides server admins with 40Gb of Ethernet bandwidth, that can be carved up into Virtual NICs (forty 1GbE ports, or more), for massive port density on a single chip.

The BCM578x0 also redefines the “Converged” in Converged Network Adapter. Starting next year, 10GbE adapter market leaders will need to support RDMA offload in order to stay market leaders. IT departments like to standardize on adapter hardware and driver images, and the BCM578x0 class of C-NICs will allow standardization on one 10GbE adapter for LAN, NAS, SAN and HPC applications. The BCM578x0 will accomplish this with hardware offload of the RDMA protocol used for low-latency networking, making it the first LOM/NIC ASIC capable of supporting offload of network protocol processing and VEB – the features needed for complete network convergence.

Looking to the future, this class of Ethernet controller with 40Gb of bandwidth has the potential to serve as a 40GbE host adapter a 40GbE iSCSI storage port, or a 40GbE FCoE storage port. In the HPC market, 40Gb (InfiniBand) server connectivity is widespread. I expect 40GbE iSCSI and FCoE storage to appear in 2012 in large data centers. The availability of 40GbE adapters will follow for direct connections to 40GbE storage and to compete with 40Gb InfiniBand for server connectivity to high performance technical and business computing applications.

- More Virtual Servers**
TCP/IP Offload
iSCSI Offload
FCoE Offload
VEB Offload
- Complete Convergence**
RDMA



- Scalable**
40Gb Available Bandwidth
4x10Gb Physical Ports
Hundreds of Virtual NICs

- Compact**
40nm technology
.82 inch square

is

BCM578x0 Family Overview

Highlights

- Quad-port, single-chip 10GbE Converged Controllers
- Industry's smallest footprint for Converged Controller: 0.82 square inches
- Bi-directional line rate on all 10G ports simultaneously
- PCI Express 3.0 support
- On-chip offload: iSCSI, FCoE, TCP, and RDMA (HW-ready)
- Virtualization: SR-IOV, NIC partitioning, Virtual Embedded Bridge (VEB)
- Power savings with EEE support for 10GBASE-T designs
- Dual-media support

Configurations

- BCM57800 dual-port 10GbE and dual-port 1GbE connectivity
- BCM57810 dual-port 10GbE
- BCM57840 quad-port 10GbE

Availability

- Sampling now, production ramp anticipated in 2H 2011



BCM578x0



Bottom Line

The bottom line is that Broadcom has raised the bar on their competition in the CNA market. To stay in the game, high-end Ethernet LOM/adaptor ASICs from Brocade, Chelsio, Emulex, Mellanox, QLogic and SolarFlare need to support 40Gb of bandwidth, at least four ports, and full offload—including VEB and some form of the RDMA protocol.

Related Links

Press Release: Multi-Port Single-Chip 10G Ethernet Converged Controllers

<http://www.broadcom.com/products/Ethernet-Controllers/Enterprise-Server/BCM578x0-Family>

BCM57810 Product Brief

<http://www.broadcom.com/collateral/pb/57810S-PB00-R.pdf>

BCM57840 Product Brief

<http://www.broadcom.com/collateral/pb/57840S-PB00-R.pdf>

About the Author



Frank Berry is founder and senior analyst for IT Brand Pulse. As former vice president of product marketing for the host bus adapter group and vice president of corporate marketing for QLogic, Mr. Berry has over 9 years experience in the development and marketing of host bus adapters. Prior to that Mr. Berry was vice president of worldwide marketing for the automated tape library (ATL) division of Quantum. You can contact Frank at frank.berry@itbrandpulse.com.