Demartek

Evaluation of IBM/QLogic FCoE Server and Storage Solution

May 2010

Introduction

IBM® and QLogic® commissioned Demartek to perform a hands-on evaluation of a server and storage solution using Fibre Channel over Ethernet (FCoE) and Fibre Channel (FC) technologies working together in an end-to-end solution. This evaluation included installing and deploying an IBM server with a QLogic Converged Network Adapter (CNA) connected to an FCoE infrastructure that was connected to an IBM DS4700 Fibre Channel storage system.

Evaluation Summary

We found that a server with an FCoE CNA connected to a DCB/FCoE switch works as expected when connected to existing native Fibre Channel storage infrastructure. From a storage management viewpoint, managing the FCoE components was equivalent to managing traditional FC components. The host operating system, applications and storage subsystem worked as expected and in the same manner as if an end-to-end Fibre Channel infrastructure had been deployed. Connecting existing native Fibre Channel storage subsystems to an FCoE infrastructure was seamless and would be straightforward for a storage administrator comfortable with Fibre Channel storage systems.

As enterprises plan new datacenters or new server and storage infrastructure, FCoE and DCB technology should be carefully examined. They offer the potential for increased performance, a reduction in the number of adapters needed and a commensurate reduction in electric power consumption, while working with existing FC infrastructure.

Overview of FCoE technology

Two Network Technologies in Today's Datacenters

Today's datacenters typically use Ethernet and Fibre Channel networks, each for its own purpose. Each of these networks uses different adapters, switches and cables today.

Ethernet networks are well-known for carrying local area network (LAN) traffic and are commonly run at 1 Gigabit per second, with 10 Gigabits per second being deployed in network cores and gradually finding deployments towards the edges of the network. Roadmaps and specifications for 40 Gb/s and 100 Gb/s Ethernet are emerging.

Fibre Channel (FC) networks are well-known for carrying storage area network (SAN) traffic and are commonly run at 4 Gigabits per second, with 8 Gigabit per second FC deployments becoming increasingly common. The roadmap for 16 Gb/s FC is well-established and 16 Gb/s FC products are expected within the next year or two.

Unified Fabrics and Datacenter Bridging (DCB)

A unified fabric is a networking fabric that combines traditional LAN and SAN traffic on the same physical network. Changes are needed in both technologies for unified fabrics.

- Ethernet must be made "lossless" and new management functions are required.
- Fibre Channel protocol must be made compatible with this new Ethernet.

The IEEE has defined the term Datacenter Bridging (DCB) to meet the requirements for this new type of unified fabric. DCB is also known by the terms Converged Enhanced Ethernet (CEE), Datacenter Ethernet (DCE) which has been trademarked by Cisco, Enhanced Ethernet for Datacenter (EEDC), and other similar terms. DCB adds four basic functions to the existing Ethernet infrastructure that enables new capabilities needed for unified fabrics:

- Traffic differentiation – DCB can distinguish between LAN, SAN and IPC traffic
- Lossless fabric – Required for SAN traffic
- Optimal bridging – Allows the shortest path bridging within the datacenter
- Configuration management – Provides configuration management functions that work with the existing infrastructures of Fibre Channel and Ethernet

There are several specifications within the IEEE 802.1 collection of standards that are expected to be ratified in 2010. These enhancements to Ethernet include selectively pausing traffic classes rather than dropping packets or pausing the entire link. This pausing function and other changes transform Ethernet into lossless Ethernet, which is required to carry storage traffic in a way that is compatible with Fibre Channel SANs. Other enhancements allow for traffic to be given differing priorities and grouping characteristics, so that administrators can guarantee a specific bandwidth and priority for differing types of traffic.

Fibre Channel over Ethernet (FCoE)

Fibre Channel over Ethernet (FCoE) is a standard that has been approved by the "T11" technical committee and works with DCB. FCoE and DCB together allow "converged" traffic to be run over the same cable simultaneously. Current deployments run the converged traffic from a unified adapter and cable in a server connected to a top-of-rack switch.

FCoE uses the existing Fibre Channel (FC) protocol running on the new lossless Ethernet (DCB). The FCoE protocol looks and acts the same as traditional FC. FCoE fabrics must be built with switches that support DCB and FCoE, and these switches must interoperate with existing FC switches, support all FC advanced features and operate identically on FC and FCoE fabrics.

Typically, the top-of-rack switches that support DCB and FCoE have varying numbers of 10-Gb Ethernet ports and optionally contain either 4-Gb or 8-Gb Fibre Channel ports. This allows these switches to handle all the LAN and SAN traffic within a rack, but forward that traffic to separate existing LAN and SAN infrastructures elsewhere in the datacenter.

Unified Adapters

A new type of "converged" or "unified" adapter is available for this new type of network. These adapters are known as converged network adapters (CNA), converged NIC (C-NIC) or other similar terms. By combining the functions of traditional FC HBAs and 10GbE NICs onto a single adapter, these adapters support FCoE and Ethernet applications simultaneously such as TCP/IP, iSCSI, etc., running at 10 Gb/s. These adapters have varying degrees of hardware offload capabilities for each of the types of traffic they support. The current generation of these unified adapters typically require servers with a minimum of x8 PCI Express (PCIe) Gen1 or x4 PCIe Gen2 slots.

Three Types of FCoE Implementations

Using a host server with a converged network adapter connected to an FCoE/DCB switch, there are three basic ways to implement FCoE with respect to the storage used.

1. Connect traditional FC storage targets to the FC ports of the FCoE switch.
2. Connect an existing FC switch to the FC ports of the FCoE switch.
3. Connect native FCoE storage targets to the 10GbE ports of the FCoE switch.

The testing conducted for this report uses option 1 above, which works with the large amount of existing native Fibre Channel storage in datacenters today.

Evaluation Environment

The evaluation was conducted at the Demartek lab in Arvada, Colorado using Demartek-owned networking infrastructure. IBM supplied the IBM server and IBM storage system used for these tests. QLogic supplied the converged network adapter (CNA).

Server

- IBM System x3650-M2
  - Dual Intel Xeon X5570, 2.93 GHz, 8 total cores, 16 logical processors
  - 16GB RAM
  - Qty. 4: x8 PCIe Gen2 slots
  - Qty. 5: 300GB, 6Gb/s SAS 2.5-inch internal disk drives
- Windows Server 2008 Enterprise x64 Edition
- IBM Storage Manager® software for DS4000™ series storage

Converged Network Adapter (CNA)

- QLogic QLE8142, dual-port 10GbE FCoE CNA

DCB/FCoE Switch

- Cisco Nexus 5020 Switch
  - Includes qty. 4: 4-Gb/s Fibre Channel ports

Storage System

- IBM System Storage DS4700® with EXP810 Storage Expansion Unit
  - Qty. 2: 4-Gb/s Fibre Channel host interfaces
  - Qty. 24: 300GB 10K RPM Fibre Channel disk drives
    (approx. 7TB of raw capacity)

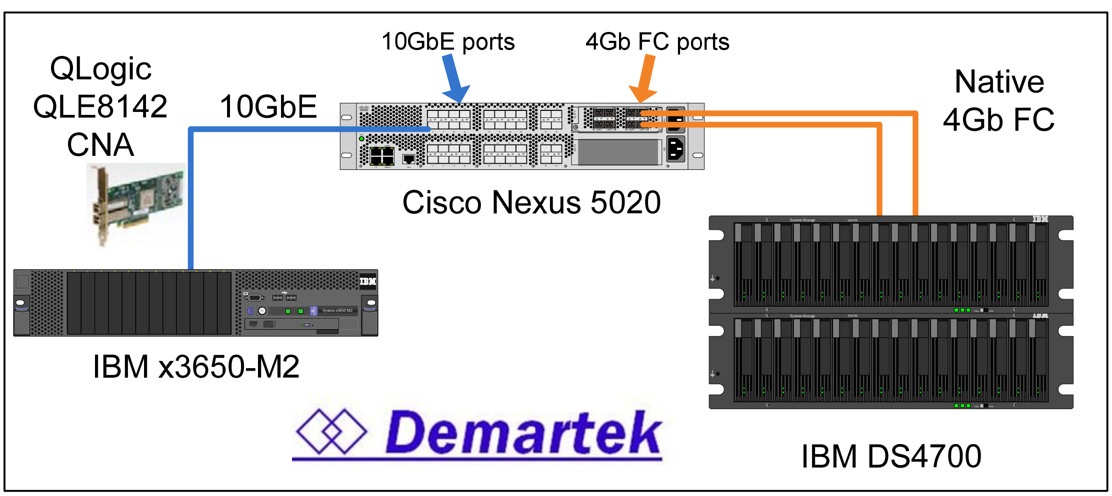

Figure 1 - Tested Hardware Configuration

Configuring the FCoE Storage Network

Configuring an FCoE storage network is much the same as configuring a native Fibre Channel storage network. Fabric zoning and LUN masking are required and the protocols and behavior of FCoE are the same as native Fibre Channel. Just like native FC HBAs, the host CNA ports require visibility to the storage. The difference is that CNAs can also run Ethernet protocols such as TCP/IP, iSCSI, etc., simultaneously with FCoE.

Switch Configuration

In the tested configuration, VLAN 300 and VSAN 3 were configured in the Cisco Nexus switch for the ports and devices shown in Figure 1. Other ports were configured for other VLANs and VSANs.

```
cisco5020# show vsan 3
vsan 3 information
         name:FibreChannel  state:active
         interoperability mode:default
         loadbalancing:src-id/dst-id/oxid
         operational state:up

cisco5020# show zoneset active vsan 3
zoneset name DMRTK-FC vsan 3
  zone name IBM-FC-Zone vsan 3
  * fcid 0x8b0000 [interface fc2/1 swwn 20:00:00:0d:ec:a3:4b:40]
     interface fc2/2 swwn 20:00:00:0d:ec:a3:4b:40
  * fcid 0x8b0100 [interface fc2/3 swwn 20:00:00:0d:ec:a3:4b:40]
     interface fc2/4 swwn 20:00:00:0d:ec:a3:4b:40
    pwwn 21:00:00:c0:dd:10:0d:79
  * fcid 0x8b0201 [pwwn 21:00:00:c0:dd:10:0d:7b]

cisco5020# show fcns database vsan 3
VSAN 3:
--------------------------------------------------------------------------
FCID        TYPE  PWWN                    (VENDOR)        FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0x8b0000    N     20:06:00:a0:b8:42:08:5d (SymBios)       scsi-fcp:target
0x8b0100    N     20:07:00:a0:b8:42:08:5d (SymBios)       scsi-fcp:target
0x8b0201    N     21:00:00:c0:dd:10:0d:7b (Qlogic)        scsi-fcp:init

Total number of entries = 3

cisco5020# show flogi database vsan 3
----------------------------------------------------------------------------
INTERFACE    VSAN    FCID           PORT NAME               NODE NAME
----------------------------------------------------------------------------
fc2/1        3     0x8b0000  20:06:00:a0:b8:42:08:5d 20:06:00:a0:b8:42:08:5c
fc2/3        3     0x8b0100  20:07:00:a0:b8:42:08:5d 20:06:00:a0:b8:42:08:5c
vfc4         3     0x8b0201  21:00:00:c0:dd:10:0d:7b 20:00:00:c0:dd:10:0d:7b

Total number of flogi = 4.
```

Figure 2 - Fabric Zoning Configuration Details

CNA Configuration

The QLogic QLE8142 CNA was installed and configured in the host server, following the standard QLogic SANsurfer® installation steps. The QLE8142 requires a server with x8 PCIe Gen1 or x4 PCIe Gen2 slots. The ports on the CNA were connected to the 10GbE ports of the Cisco Nexus switch. The tests were run with a single CNA port connected and with both CNA ports connected.

QLogic's SANsurfer behaved the same way that it does for managing the QLogic Fibre Channel adapters, except for the addition of Ethernet-specific configuration data that is available with CNAs.

For the initial configuration, one port of the CNA was configured into the same VSAN (VSAN 3) as the IBM storage system. The other port was configured into a different VSAN.

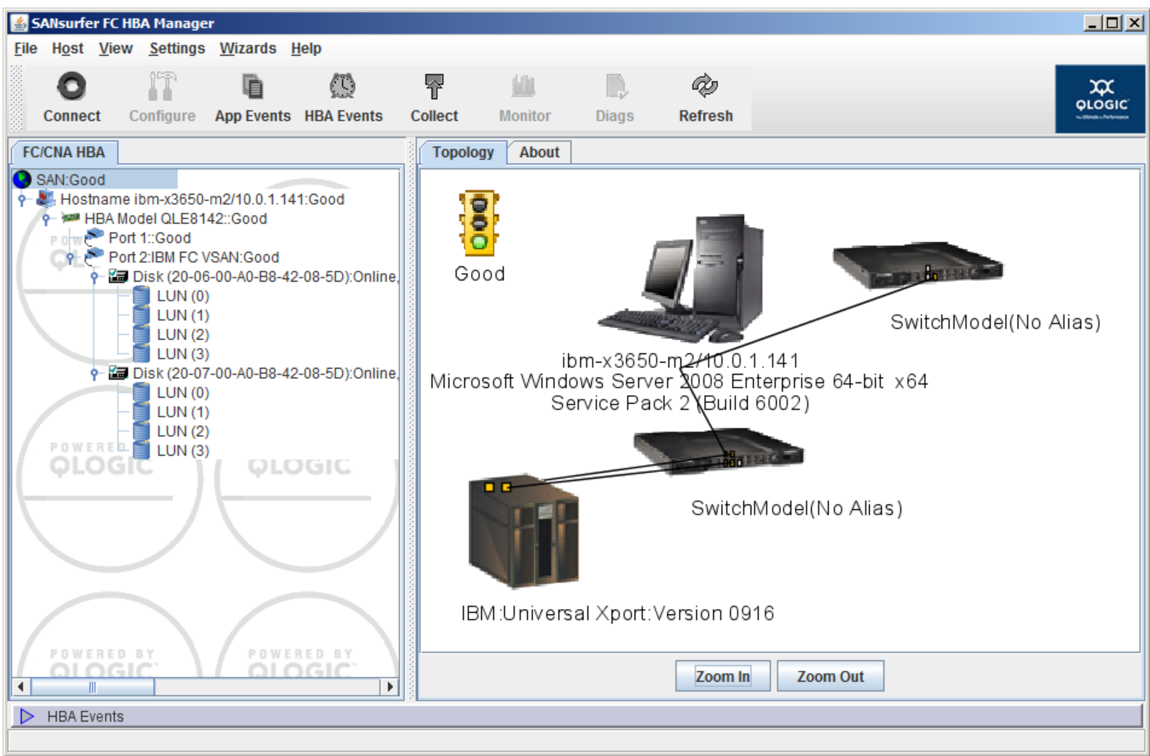

Figure 3 - QLogic SANsurfer configuration

Storage Configuration

The two native 4 Gb/s Fibre Channel host ports of the IBM DS4700 were connected to the 4 Gb/s Fibre Channel ports of the Cisco Nexus switch. Because the host server with the CNA was also used as the storage management server, the storage management software was able to manage the storage using in-band and out-of-band connections.

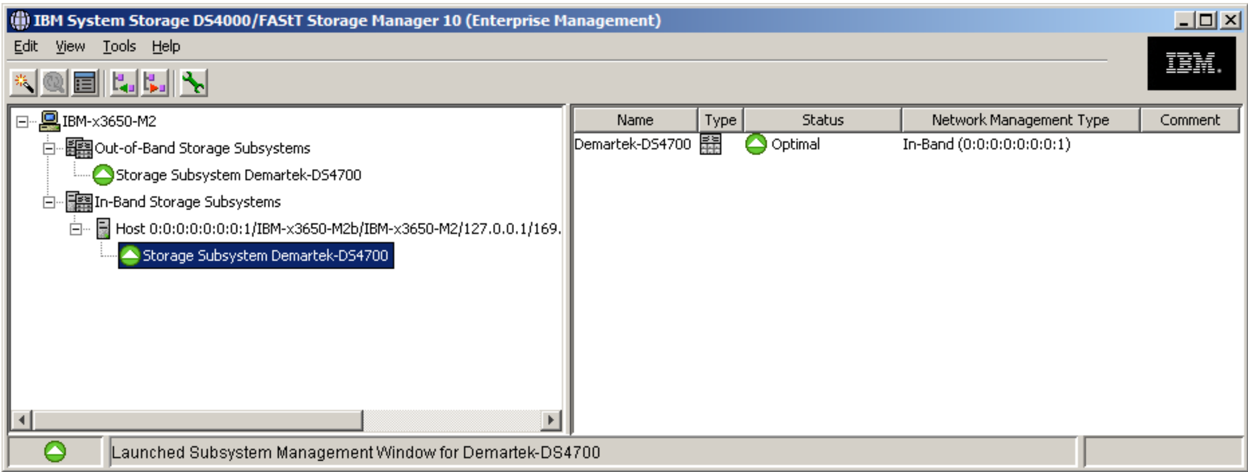

Figure 4 - IBM Storage Manager

Three logical drives (LUNs) were created and assigned to the IBM x3650-M2 host server. The logical drives were created as RAID10 sets.

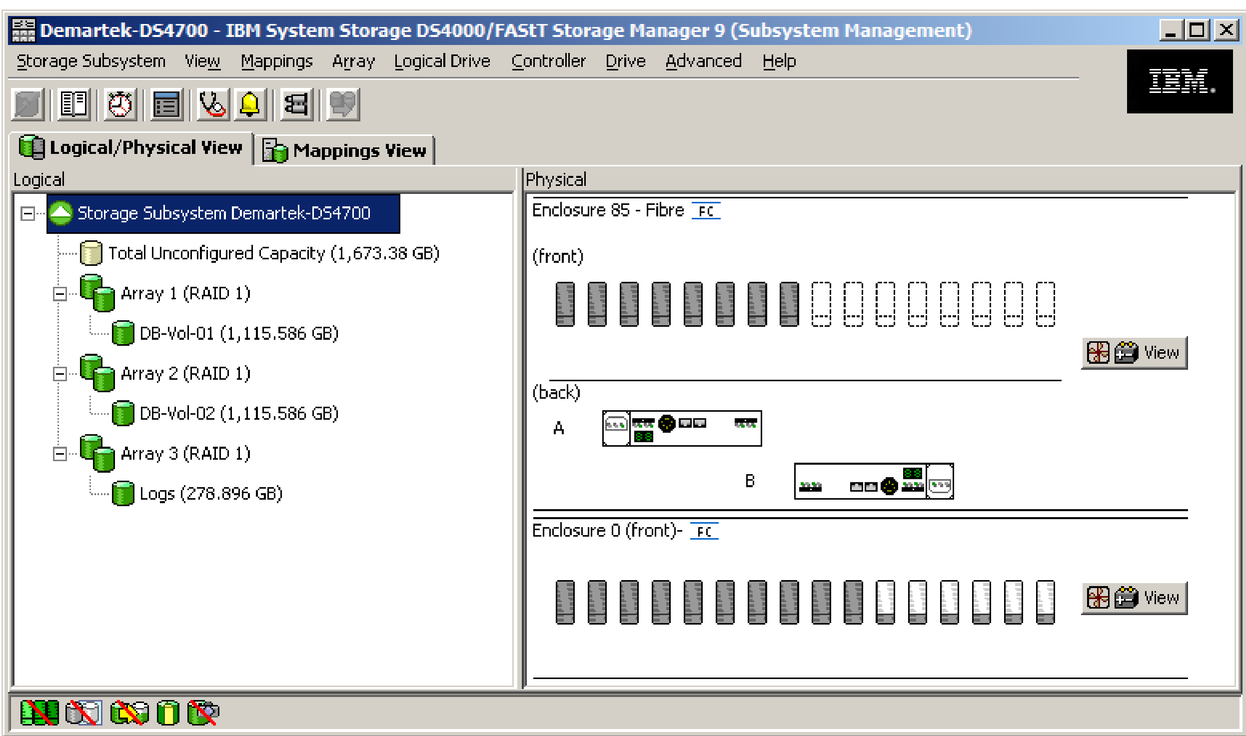

Figure 5 - IBM Storage Manager Details

Host Configuration

After configuring the switch and storage, the server viewed the storage exactly the same way that it does in a native Fibre Channel environment, and the storage viewed the host server exactly the same way that it does in a native Fibre Channel environment.

The IBM Storage Manager software includes management functions for the storage array and the multi-path support for the storage array. The host operating system correctly saw only one instance of each LUN even though there were multiple paths to the storage system.

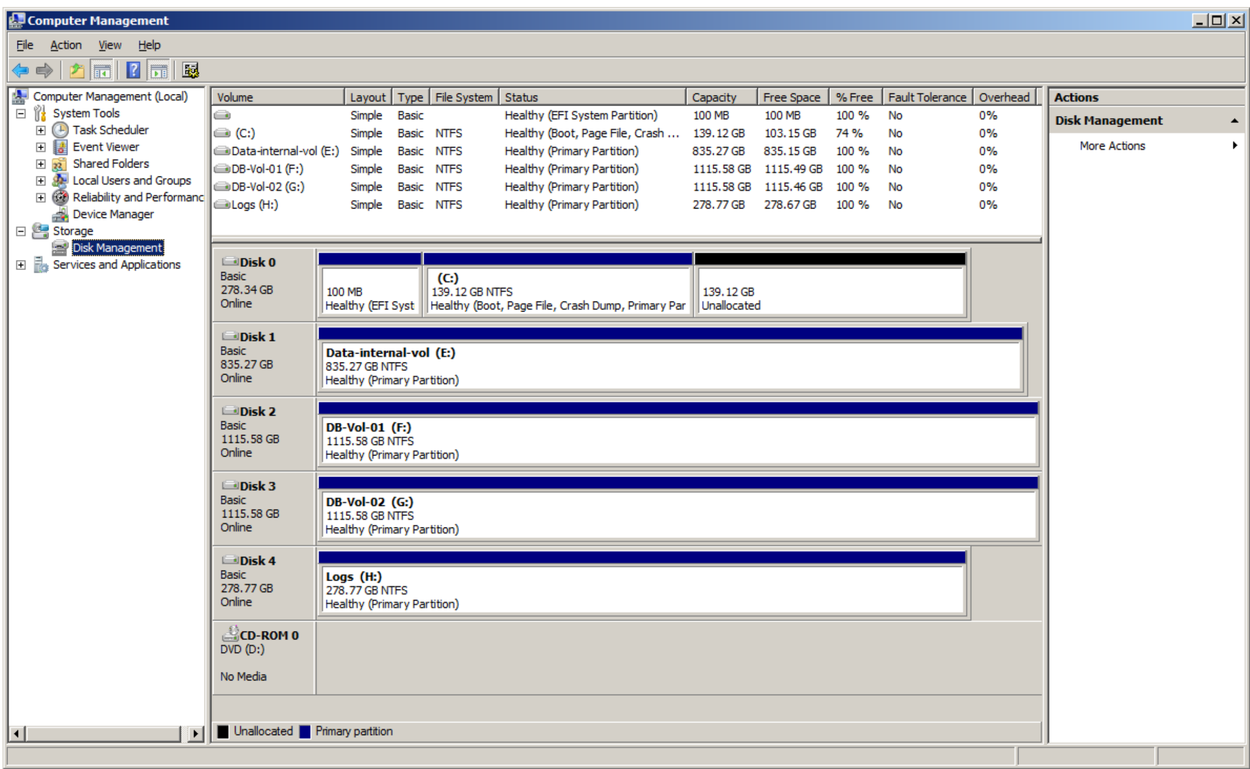

Figure 6 - Host Server Disk Manager

The host applications were able to use the storage in the usual manner.

Test Results

After configuring the FCoE storage network, the server, CNA and storage behaved and performed exactly as expected. The storage performance was exactly as expected for this type of storage with two native 4 Gb/s Fibre Channel host ports and two disk shelves of 300GB 10K RPM disk drives. The storage traffic flowed from the host server through the 10GbE ports (1 or 2) of the CNA to the 10GbE ports on Cisco Nexus DCB/FCoE switch. The traffic was sent through the 4Gb FC ports of the Cisco switch to the 4Gb FC host ports of the storage system.

The host operating system and applications viewed the storage as expected and the applications had no knowledge of the fact that they were using an FCoE connection to the storage. LUNs could be created on the storage system and made visible to the host operating system in the usual manner.

Conclusion

We found that FCoE technology makes it possible to deploy a converged network adapter with today's generation of new servers that works seamlessly with existing Fibre Channel storage systems. The deployment and management of the converged network adapters in new servers is almost identical to the deployment and management of traditional Fibre Channel adapters in similar servers.

From a storage management perspective, the installation and management of the FCoE components is very similar to the equivalent installation and management of native FC infrastructure.

As new server infrastructures are deployed, whether in new datacenters or for new server deployments in existing datacenters, FCoE components provide the potential to reduce the total number of adapters and cables required, lowering the power consumption and heating load while increasing the total I/O bandwidth available to new servers. In addition, FCoE technology rides on the roadmap of the new Ethernet, opening up the possibility for significant performance improvements as Ethernet moves to 40 Gb/s and 100 Gb/s in the future.
