EMULEX

TECHNICAL BRIEF

CONNECTIVITY

Next Generation iSCSI Solutions from Emulex and Dell

iSCSI storage has typically been viewed as the "low-cost" SAN option—good for small and medium sized businesses or branch office deployments. That profile is changing. Emulex and Dell have teamed together to deliver an end-to-end iSCSI solution based on Data Center Bridging (DCB) standards developed by IEEE. With Emulex and Dell, you have the products and support that you need for a truly enterprise-class deployment.

iSCSI over DCB is an end-to-end Dell and Emulex solution that delivers centralized configuration, management and server integration while enabling deterministic performance for converged LAN and SAN traffic.

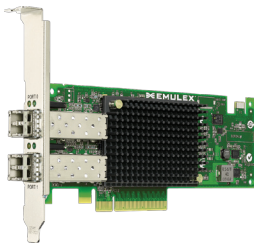

Emulex OCe11102-I iSCSI CNA

Dell PowerConnect 10GbE Switch

Dell EqualLogic iSCSI Storage Array

Features and Benefits

10Gb Ethernet

OneConnect OCe11102-I Converged Network Adapters (CNAs) are dual-port 10Gb Ethernet (10GbE) adapters that resolve the limitations with the current generation of 1Gb NICs and iSCSI HBAs. OneConnect OCe11102-I CNAs provide the bandwidth needed for virtualized servers and applications such as databases and streaming video that demand high I/O rates.

Data Center Bridging

One of the challenges with traditional iSCSI deployments is the non-deterministic nature of Ethernet. Packets can be lost, requiring retransmission of data and unacceptable I/O delays. There's also no capability to assign priorities to classes of I/O. As a result, iSCSI best practices require separate Ethernet networks for iSCSI traffic. This means dedicated NICs for the server and separate switches or dedicated Virtual LANs (VLANs) in the network.

DCB standards provide the framework for high-performance iSCSI deployments with key capabilities that include:

- **Priority Flow Control (PFC)**—enables "lossless Ethernet", a consistent stream of data between servers and storage arrays.
- **Quality of Service (QoS) and Enhanced Transmission Selection (ETS)**—support protocol priorities and allocation of bandwidth for iSCSI and IP traffic.
- **Data Center Bridging Capabilities eXchange (DCBX)**—enables automatic network-based configuration of key network and iSCSI parameters.

With DCB, data centers can use a single set of 10GbE adapters and switches for storage and network traffic. PFC eliminates retransmission of data and supports a consistent flow of data with low latency. With QoS and ETS, iSCSI can be assigned the bandwidth and priority that's required to support critical applications.

Based on three generations of Emulex CNA technology, OCe11102-I iSCSI adapters enable Dell PowerEdge servers to connect to high-performance DCB-based networks.

DELL

Next Generation iSCSI Solutions from Emulex and Dell

iSCSI Offload

OCe11102-I CNAs support full iSCSI protocol offload, providing performance that is superior to a NIC and iSCSI software initiator. Critical CPU resources are used to enable more virtual machines (VMs) per physical server and support demanding applications.

With iSCSI offload, a separate storage device is presented to the operating system or hypervisor and a separate driver stack is used to process storage I/O. In contrast, a software iSCSI initiator runs on top of the TCP/IP stack so storage and network traffic use the same processes. The result is limited scalability with the NIC and software initiator option.

More Virtual Machines per Server

Emulex Labs conducted a series of benchmark tests to evaluate VM scalability comparing the OCe11102-I iSCSI adapter with software ISCSI using a 10GbE NIC. The goal was to compare the maximum number of VMs that could run concurrently with a constant I/O rate. The maximum number of VMs was reached when the I/O throughput dropped below the constant rate.

Two combinations of I/O rates and block sizes were used:

- 10,000 IOPS (I/Os per second) with 4Kb block size
- 5,000 IOPS with 8Kb block size

The following results were observed:

- 58% more VMs with 4Kb block size for the OCe11102-I
- 53% more VMs with 8b block size for the OCe11102-I
- 50% lower watts per VM for the OCe11102-I

As shown with these tests, protocol offload is a critical technology to maximize the benefits of server virtualization.

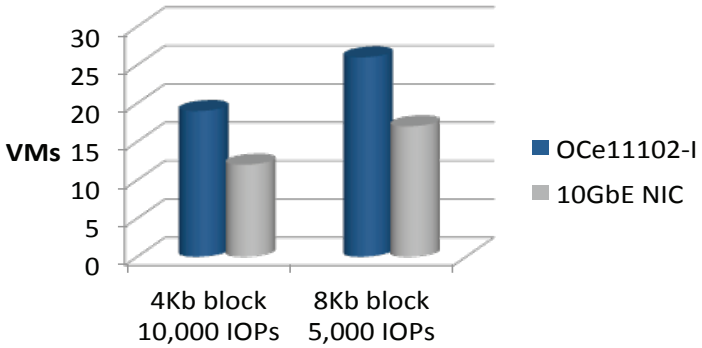

Figure 1 Concurrent VMs with constant I/O rates.

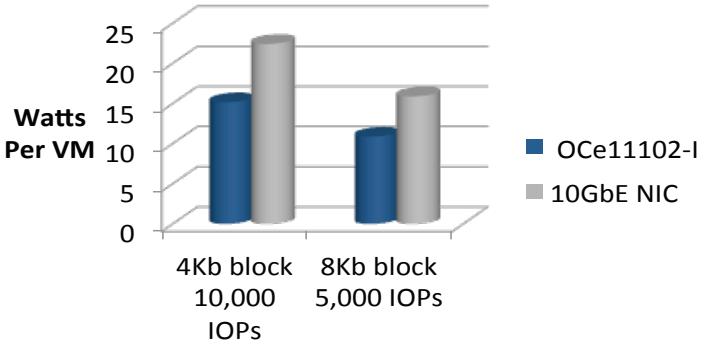

Figure 2 Server watts per VM.

Boot from SAN

OCe11102-I CNAs support boot from SAN using an iSCSI hardware device. Boot targets and parameters are easily set using boot management software that runs on adapter, making the OCe11102-I an ideal solution for diskless deployments.

Universal Multi-Channel

OCe11102-I CNAs also support Universal Multi-Channel (UMC), which enables each physical port to present multiple PCI functions to the operating system or hypervisor. This gives further options to segment and manage bandwidth for network and storage I/O. Each port of the OCe1102-I adapter provides one iSCSI function and three NIC functions that can be managed as unique physical devices.

OneCommand™ Manager

OneCommand Manager provides centralized management of Emulex One Connect adapters and LightPulse® HBAs throughout the data center from a centralized management console. Centralized discovery, monitoring and management of local and remote adapters can be done from a secure remote client. In-depth management capabilities include firmware and boot code upgrades, beaconing, statistics and advanced diagnostics.

World Headquarters 3333 Susan Street, Costa Mesa, CA 92626 +1 714 662 5600
Wokingham, UK +44 (0) 118 977 2929 | **Munich, Germany** +49 (0) 89 97007 177
Paris, France +33 (0) 158 580 022 | **Beijing, China** +86 10 68499547
Tokyo, Japan +81 3 5325 3261 | **Bangalore, India** +91 80 40156789

Connect with Emulex twitter.com/elx4dell friendfeed.com/emulex bit.ly/emulexfb bit.ly/emulexlinks bit.ly/elxdellblog

www.emulex-dell.com **www.emulex.com/dellblog**

11-1145 · 6/11