



# Virtual Desktop Infrastructure (VDI) Acceleration Using LSI and Mellanox Low Latency Interconnect Technologies



Every IT professional's goal is to improve TCO and user experience. In a Virtual Desktop Infrastructure (VDI) application, the objective is to increase the efficiency by supporting the maximum number of desktops per server while maintaining response times to users that would be comparable to a physical desktop. In addition the solution must be resilient since downtime of the VDI application causes the idling of hundreds to thousands of users and consequently reduces overall organizational productivity and increases user frustration.

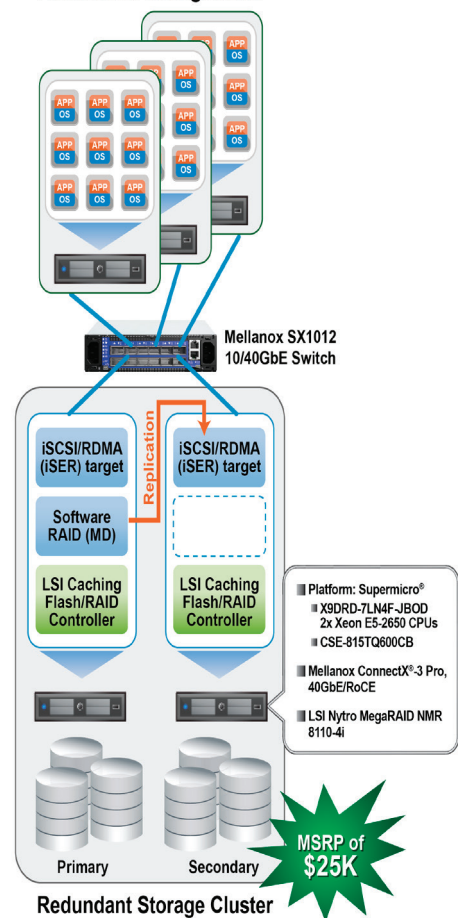
Legacy Fibre Channel-connected storage subsystems provide shared storage which enables moving virtual machines between physical servers. Unfortunately, FC transport cannot keep up with today's fast storage options and solid state drives (SSD). Leveraging an existing Ethernet infrastructure saves costs by combining networking and storage I/O over the same cable. Deploying SSDs in the subsystem can support the challenging I/O demands from a VDI environment, but completely replacing HDDs is expensive and often unnecessary. Tiered solutions that combine both SSD and HDD can provide the best of both worlds.

## Virtual Desktop Infrastructure (VDI)

VDI enables companies to centralize all of their desktop services over a virtualized data center. With VDI users are not tied to a specific PC and can access their desktop and running applications over their data from any location. VDI also helps IT administrators to have more efficient and secure environment which enables them to serve better their client business needs.

VDI efficiency is measured by the number of virtual desktops that a specific infrastructure can support within less than a specific time. The lim-

## Benchmark Configuration



**Figure 1. LSI/Mellanox Storage Accelerator**

iting factor is the access time to storage. Using smart disk controller with caching mechanism combined with fast interconnect that support RDMA significantly eliminates potential bottlenecks which enables the lowest Total Cost of Ownership (TCO) and highest efficiency.

## Low Latency Solutions

Low latency data requests from storage or other servers are the key to enabling more VDI sessions without increasing user response times. LSI Corporation and Mellanox Technologies are leaders in low latency storage and networking solutions and have combined LSI storage solutions with Mellanox high-performance networking to provide a Storage Accelerator. LSI's Nytro MegaRAID application acceleration card provides flash storage using a server's PCIe backplane. It uses an economical mix of flash and HDD storage and uses adaptive caching algorithms to cache heavily referenced areas of storage. Mellanox's 40GbE with Remote Direct Memory Access (RDMA) technology provides high bandwidth and enables the fastest and most reliable data transfers with minimal server compute resources. By using the iSCSI protocol with an RDMA transport (a standard called iSER) instead of iSCSI over TCP/IP, it leads to significantly faster storage latency and bandwidth while maintaining the familiarity of iSCSI management and utilities in the VMware vSphere environment.

## RDMA over Converged Ethernet (RoCE)

RoCE utilizes advances in Ethernet to enable efficient implementations of RDMA over Ethernet. It enables widespread deployment of RDMA technologies in mainstream datacenter applications. RoCE-based network management is the same as that for any Ethernet network management, eliminating the need for IT managers to learn new technologies.

The business advantages of RoCE for a data center are:

- No changes to data center infrastructure
- I/O unification on a single wire over 10/40 GbE networks
- Continuation of existing data center management infrastructure
- Reduction in power and cost savings
- Maintain existing and future application compatibility
- Significant CapEx and OpEx savings with a single chip solution for I/O unification

## iSCSI Extensions for RDMA (iSER)

iSCSI Extensions for RDMA (iSER) is an IETF standard extension to the Internet Small Computer System Interface (iSCSI) protocol. It uses hardware Remote Direct Memory Access (RDMA) for the iSCSI transaction delivery mechanism. It basically uses the upper layers of iSCSI for session management, discovery, recovery, and so on, and is therefore compatible with all the features and func-

tions supported by iSCSI. However, using iSER eliminates the bottleneck and provides unprecedented performance through the following mechanisms:

- Eliminates data copying via RDMA technology
- CRC is calculated by hardware to reduce CPU overhead while maintaining data integrity
- Works with message boundaries instead of streams, allowing faster protocol processing
- Transport protocol is fully implemented in hardware (requiring minimal CPU cycles per IO and driving extremely low latency)

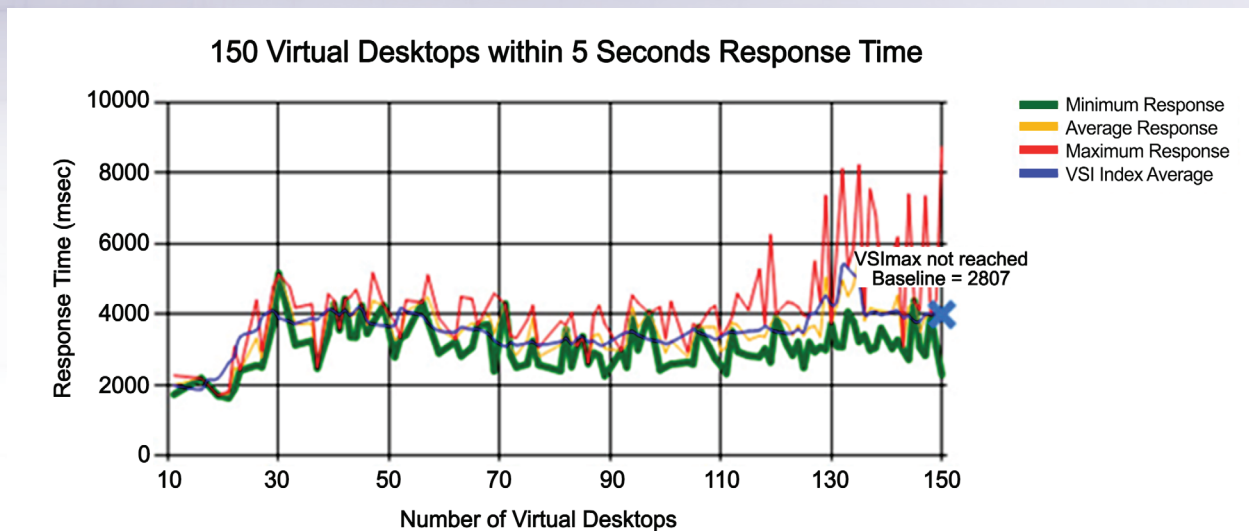
iSER implementations today can drive magnitudes faster latency and bandwidth than traditional FC or iSCSI implementations.

Mellanox iSER driver for vSphere has been certified under VMware's Partner Verified and Supported Products (PVSP) program and can be found at: <http://www.vmware.com/resources/compatibility/vcl/partnersupport.php>

## High Performance Can Be Economical

To evaluate VDI performance using the LSI/Mellanox Storage Accelerator (see [Figure 1](#)), a Supermicro X9DRD-7LN4F-JBOD serverboard in a 1U Supermicro CSE-815TQ-600CB chassis with a Xeon E5-2650 processor was used running a Linux operating system. A Mellanox SX1012 12-port 40GbE switch in a half-width 1U form factor allows multiple ESXi servers using iSER to communicate with the Storage Target (which is based on an open source Linux iSCSI and iSER target implementation). Mellanox ConnectX®-3 Pro adapter supports 40Gb/s RoCE storage connectivity between the servers. Providing data storage is a Nytro MegaRAID NMR 8110-4i card which has 200GB of on-card flash and eight SAS HDDs.

To facilitate continuous VDI operation, a secondary server retains a copy of the storage on the primary server using software mirroring across another iSER connection. If the primary server fails, the secondary server supports storage requests. Altogether, the MSRP for the hardware used in this evaluation was about \$25,000 for a fault-tolerant, highly available, flash storage system with high performance 40GbE interconnects.



**Figure 2. VDI Performance**

### Evaluating VDI Performance

VDI performance was measured using Login VSI's VDI load generator which creates the actual workload of a typical Windows user using Microsoft Office applications. Figure 2 shows 150 virtual desktop sessions running in parallel while each user experiences less than 5 seconds in response time. Booting all the 150 sessions was accomplished in less than 6 minutes.

These results represent more than 2 times the efficiency boost versus 10GbE and iSCSI over TCP/IP, in which only 60 virtual desktop sessions could run with 5 seconds response time using the same setup.

### Scaling for Additional Desktops

The Storage Target can easily scale the number of virtual desktops. It can handle increments of 150 virtual desktops connected to an vSphere server. Deriving from measured key metrics, the Storage Target can support up to 1500 virtual desktops using two Nytro MegaRAID cards. More desktops can be supported by adding Nytro MegaRAID cards.

### Failover

Failover functionality was not implemented in this performance benchmarking. However, open-source implementations like Linux HA and Linux DRBD® have failure detection and IP address migration, and could be used to add and test this capability.

## KEY SOLUTION ADVANTAGES

- Maximizing number of virtual desktop per server with guaranteed response time
- Most resilient scale-out solution minimizing downtime of business applications running over VDI
- Improves data center efficiency by enabling higher performance and a lower cost of ownership
- Consolidates LAN and SAN infrastructures into a single Mellanox high-performance interconnect infrastructure of Ethernet and/or InfiniBand adapters, cable, and switches



350 Oakmead Parkway, Suite 100, Sunnyvale, CA 94085  
 Tel: 408-970-3400 • Fax: 408-970-3403  
[www.mellanox.com](http://www.mellanox.com)